The University After Generative AI

A Field Guide for Designing AI-Native Higher Education, 2026–2035

SILVIO MEIRA • LUCIANO MEIRA

tds.company

The University After Generative AI

A Field Guide for Designing AI-Native Higher Education, 2026–2035

SILVIO MEIRA[1,2,3] AND LUCIANO MEIRA[2,4]

Version 6 · July 2026

ISBN 978-65-978626-4-1

About the Authors

Silvio Meira was born in Taperoá, Paraíba, in 1955. He trained as an electronic engineer at the Instituto Tecnológico de Aeronáutica (ITA, 1977), took a master's degree in informatics at the Universidade Federal de Pernambuco (UFPE, 1981), and completed a PhD in computing at the University of Kent at Canterbury, in the United Kingdom (1985). He is a founder and Chief Scientist of TDS.company, Extraordinary Professor at CESAR School, Distinguished Research Fellow at the Asia School of Business in Kuala Lumpur, and Professor Emeritus at the Centro de Informática of UFPE. He is a founder of Porto Digital, where he chairs the board, and in 1996 he founded CESAR — the Recife Center for Advanced Studies and Systems — which he led until 2014. A former fellow of the Berkman Klein Center at Harvard University (2012–2015), he has authored more than 300 scientific papers and over a thousand texts on information technology and its social impact, including the books O Que É Estratégia? (2023) and Everything Platforms (2025), and has supervised more than 150 doctoral and master's theses. His honours include the Ordem Nacional do Mérito Científico (1999) and Brazil's Eminent Engineer (2023).

Luciano Meira was born in Cajazeiras, Paraíba, in 1962. He is a developmental and cognitive psychologist whose work sits at the intersection of learning, culture and educational technology. He has retired as an associate professor of psychology at the Universidade Federal de Pernambuco (UFPE), and is now a research-professor at CESAR School, and Head of Pedagogy at Proz Educação. He holds a PhD in mathematics education from the University of California at Berkeley (1991), a master's degree in cognitive psychology from UFPE (1987), and a bachelor's degree in pedagogy from FAFIRE (1984). Since 2017 he has been a Lemann Fellow and Visiting Scholar at the Lemann Center for Educational Entrepreneurship and Innovation at Stanford University. His research addresses human development, sociocultural and socio-historical theory, sense-making and the use of digital technologies in learning. A recognised specialist in gamification, he has led educational-innovation ventures — including as a co-founder at Joy Street — that use games and playful technologies to strengthen basic education, and has written widely on cognitive psychology, culture, development and learning.

Silvio and Luciano Meira are brothers —
a lifelong dialogue between computing and psychology that this paper continues.

Contents

Abstract

Generative and agentic artificial intelligence are **not** merely new educational technologies. They are general-purpose infrastructures of informaticity — the historically realised space of computation, communication and control whose scaffolding was laid by **Turing** (1936), **Shannon** (1948) and **Wiener** (1948) and whose contemporary architecture is analysed in **Meira** and **Meira** (2026) and **Meira**, **Neves** and **Braga** (2025). As infrastructure, AI alters the scarcity of explanation, feedback, translation, programming, assessment, drafting and analysis. Its effect on higher education is not a tool problem; it is a redesign problem at the level of institutional architecture.

This paper develops a policy review and conceptual framework for higher education in 2026–2035. Its central argument is that higher education must be understood as a sociotechnical hyperobject — a term borrowed, with deliberate qualifications, from **Morton** (2013) — and reconstructed around AI-native principles. It distinguishes AI-enabled higher education, which adds AI tools to inherited architectures, from AI-native higher education, which redesigns learning, assessment, credentials, data governance, academic work, business models and external partnerships around ubiquitous AI as cognitive infrastructure. The question is no longer how institutions adopt AI, but how they educate, evaluate, certify and organise academic work when AI is everywhere.

The framework is a critical policy synthesis, not a systematic review, and it separates three levels of claim — robust claims, plausible propositions, and explicitly normative design proposals with their refusal conditions. It draws on recent labour-market evidence (**Eloundou** et al., 2024; **Brynjolfsson**, **Li** and **Raymond**, 2025; **Brynjolfsson**, **Chandar** and **Chen**, 2025; **Hampole** et al., 2025; **Autor** and **Thompson**, 2025; **Massenkoff** and **McCrory**, 2026; **Acemoglu** and **Restrepo**, 2019), on research on assessment and evaluative judgement (**Boud** and **Falchikov**, 2006; **Bearman** et al., 2024), on the institutional theory of the university (**Trow**, 1973; **Kerr**, 1963; **Clark**, 1983; **Barnett**, 2011; **Marginson**, 2016; **Slaughter** and **Rhoades**, 2004, **Fugetta**, 2012), and on current policy frameworks (**UNESCO**, 2026; **OECD**, 2026; **Council** of the **European Union**, 2022; **European Parliament** and **Council**, 2024; **NIST**, 2023, 2024).

The argument is situated in the scale and fragility of a planetary system. Global enrolment surpassed 269 million students in 2024 (**UNESCO**, 2026), yet expansion has not produced equalisation: the stratification **Trow** (1973) anticipated persists, while demographic decline, fiscal pressure and eroding trust destabilise the inherited model (**Clark** et al., 2026; **UPCEA**, 2025). AI intensifies the unbundling of the historic bundle — instruction, certification, formation, sociality and credentialing — and the platform dynamics analysed in **Meira** (2025a) show how value then migrates to whoever controls the orchestration layer.

The defensible value of the university lies not in content delivery but in the formation and credible certification of capabilities that require judgement, accountability, social learning and verified human performance.

The most immediate pressure falls on the graduate entry contract. Because entry-level work has consisted largely of the routinised tasks AI now performs, the apprenticeship rungs through which novices became experts are removed even as demand for expert judgement rises — an apprenticeship paradox that echoes **Polanyi** (1966) on tacit knowledge and the expertise dynamics formalised by **Autor** and **Thompson** (2025). The same paradox operates inside the institution, where AI can relocate academic labour toward mentoring or fragment it into dashboard-mediated digital Taylorism (Meira, Neves and Braga, 2025). That students already use AI for work still assessed as individual (**Stephenson** and **Armstrong**, 2026) makes the equity stakes concrete: without deliberate design, AI widens the gap between a prestige bundle and an access bundle.

In response, the paper specifies nine constitutive features of an AI-native institution, grounded in the triadic intelligence ontology of **Meira** et al. (2026): individual intelligence (embodied cognition, emotion and creativity), social intelligence (the coordinating capacity of networks and institutions) and artificial intelligence (functional agency without a phenomenological substrate), whose super-additive composition is the substance of meta-cognitive work. From this follow a five-layer curriculum that protects foundational capability while distributing AI and data fluency, and an assessment regime built on authentic performance, process visibility and live defence rather than the unsupervised written artefact (**Boud** and **Falchikov**, 2006; **Bearman** et al., 2024).

The hyperobject framing is developed into fourteen aligned dimensions — from purpose and epistemic order to academic work, ecosystems and public value — used not as a taxonomy but as a diagnostic tool for whether a reform moves the dimensions that must move together. Beneath them lies the informaticity substrate formalised in **Meira** and **Meira** (2026): the axes of computability, communicability and controlability, and the triaxial phygital ontology of the physical, the networked-digital and the networked-social. The transition is read through the red-swan register of **Meira** (2025b), which names ruptures whose ontology inherited strategic planning cannot accommodate.

To bridge concept and practice, eight design vignettes — from a biology student and a history seminar to a medical clerkship, a legal-writing studio, a mobile-first vocational pathway and a Global-South community hub — render AI-native composition as daily institutional life. Reform is then located in the composition of institutions, organisations and ecosystems, adapting **North** (1990) to explain why single-layer reforms fail. A four-horizon implementation architecture (ninety days, twelve months, three years and ten

years) and a seven-level maturity model provide a navigable pathway, answering the provost's Monday-morning question of where to begin.

At the level of policy, the paper proposes ten principles and a set of instruments — recognition of modular credentials (**Council** of the **European Union**, 2022), learner data rights and risk-based AI governance (**European Parliament** and **Council**, 2024; **NIST**, 2023, 2024), public-interest procurement and revised licensure — while warning that outcomes-based assurance can reproduce the metric-gaming and capture it is meant to prevent. Transformation is a political process, not an optimisation: drawing on the Schumpeterian image of creative destruction and the academic-capitalism analysis of **Slaughter** and **Rhoades** (2004), it asks who bears the destruction and who captures the creation, and treats faculty, students, staff and unions as co-designers.

The Global South is treated as a site of theory rather than of deficit or adoption, foregrounding infrastructure, multilingualism, academic labour and data sovereignty as design questions whose relevance is ultimately universal. The paper specifies seven risks — platform capture, stratification, cognitive outsourcing, assessment collapse, faculty degradation and digital Taylorism, regulatory capture and mission drift — and defends refusal as a governance capacity that distinguishes universities from the platforms that imitate them. Five scenarios for 2035 and ten comparative research priorities with indicators turn AI-native higher education into a testable object of study rather than a slogan.

The argument is finally reflexive: a university cannot credibly prepare students for meta-cognitive orchestration while organising its own labour as programmed compliance. The creative destruction AI demands must occur within the institution, not only around it. The university that deserves to exist after generative AI is neither a nostalgic campus defended against machines nor a platform that automates learning while calling itself education, but **a human institution redesigned for a world of technologically mediated cognition** — one whose purpose is to keep the people who use these systems capable of judgement, accountability and public reason.

Keywords: higher education; generative AI; AI-native education; sociotechnical hyperobject; assessment; credentials; academic labour; triadic intelligence; meta-cognitive orchestration; informaticity; phygital space; red swans; Global South; political economy of change.

1. Introduction: From Tool Adoption to Institutional Reconstruction

1.1 The Compositional Crisis of the University

The university has survived printing, industrialisation, state-building, mass schooling, broadcast media, computing, the internet, platformisation and successive waves of educational technology. Survival, however, is not the same as continuity of form. Across nine centuries the institution has repeatedly redesigned itself: the medieval corporation of masters and scholars, the Humboldtian research university, the American multiversity that **Kerr** (1963) described as a precarious truce among incompatible communities, the mass and universal systems that **Trow** (1973) tracked as enrolment ratios climbed past the elite threshold, the entrepreneurial and corporatised university analysed by **Slaughter** and **Rhoades** (2004) under the rubric of academic capitalism, the "common good" framing recovered by **Marginson** (2016), and the disciplined defence of *being a university* offered by **Barnett** (2011). Each redesign coexisted with the previous one rather than replacing it. The contemporary university is a palimpsest of these forms — a fact that makes its present crisis hard to perceive as a single problem.

The historically dominant architecture of higher education — fixed cohorts, time-bounded degrees, lecture-and-examination teaching, faculty-mediated authority, paper credentials and place-based campuses — is now stressed simultaneously along several axes. Demographics, fiscal capacity, public trust, geopolitical competition and the unbundling of the labour-market signal are each independently pressuring the inherited form (Clark et al., 2026; UPCEA, 2025; **UNESCO**, 2026). Generative artificial intelligence is the latest of these stressors and, in some respects, the most disruptive, because it does not arrive as a single technology. It arrives as infrastructure.

Following the discussion about metacognitive work developed in **Meira**, **Neves** and **Braga** (2025) and the ontology of phygital space proposed in **Meira** and **Meira** (2026), generative AI must be read as the latest historical realisation of the triple axis of computability, communicability and controlability that constitutes contemporary informaticity. The three axes were sketched at mid-twentieth-century by the foundational papers of **Turing** (1936) on universal computation, **Shannon** (1948) on the mathematical theory of communication, and **Wiener** (1948) on cybernetics; their convergence into a single deployable infrastructure is the historical event that distinguishes the present moment from earlier waves of educational technology. Treated as infrastructure, AI changes the conditions under which universities can **justify** their existence — not because it makes them obsolete,

but because it changes the cost and availability of operations on which their inherited model silently relied.

This paper develops a policy review and conceptual framework for higher education in the period 2026–2035. Its central claim is that AI transforms higher education **not by** adding one more device to classrooms **but by** altering the cost and availability of cognitive operations on which the inherited model depended: *explanation, feedback, dialogue, translation, drafting, programming, summarisation, administration and routine assessment*.

> When the marginal cost of these operations approaches zero, the **bundle** that universities have historically sold — *instruction, certification, formation, sociality and credentialing* — comes apart unless it is consciously redesigned.

The compositional crisis is not, therefore, a technology problem to be solved by procurement. It is a redesign problem at the level of institutional architecture, governed by the same logic that **Meira** (2025a) applies to incumbents in other sectors when their value chain is disintermediated by a general-purpose digital infrastructure.

1.2 Why "AI-Enabled" Is Not Enough

Most institutions today are (somehow) AI-enabled. They have written acceptable-use policies, licensed copilots, piloted tutors, and begun staff training. These steps are necessary, but they are not sufficient, and they often confuse motion with direction.

> An AI-enabled institution still organises learning, assessment, credentials, data governance, faculty work and business models around assumptions that AI silently invalidates — **that** text production is a reliable proxy for thinking, **that** the unsupervised written assignment certifies the individual student, **that** lectures are the principal mechanism of explanation, **that** faculty effort distributes itself optimally across teaching, research and service, **that** credentials accumulate inside the institution rather than across institutions and providers, **that** learner data are mainly a compliance matter, and **that** the institutional mission can be served by adopting tools without re-examining the architecture into which they are inserted.

Each of these assumptions is now **contestable**. Bearman and colleagues (2024) document, in some detail, how generative AI changes the conditions under which evaluative judgement can be developed and assessed in higher education. **Boud** and **Falchikov** (2006), writing well before the generative AI moment, established the underlying principle: *assessment that does not align with long-term learning will eventually be gamed, mis-signalled or hollowed out*. **Stephenson** and **Armstrong** (2026), reporting on the 2026 HEPI student survey, show that students are routinely using generative AI for tasks that

universities still treat as authentic individual work. Selwyn-style (**Selwyn**, 2019) scepticism about the educational claims of AI is warranted; so too is the recognition that the existing assessment apparatus has been broken for longer than the AI debate acknowledges.

The deeper move is therefore **not a** faster adoption of AI tools **but a** redesign of the underlying logic of the institution around the existence of ubiquitous AI as cognitive infrastructure.

> The question is no longer "how do we adopt AI?"
>
> It is "how do we educate, evaluate, credential and organise academic work when AI is everywhere?"

That question requires a **different** vocabulary, a **different** governance posture and a **different** relationship with the institution's own labour — the relationship that **Meira**, **Neves** and **Braga** (2025) analyse as the transition from programmed labour to meta-cognitive orchestration.

1.3 What This Paper Does, and What It Does Not

The paper makes six contributions. **First**, it sharpens the distinction between AI-enabled and AI-native higher education and specifies nine constitutive features of an AI-native institution. **Second**, it treats higher education as a sociotechnical hyperobject and identifies fourteen aligned dimensions, with explicit acknowledgement of the limits of the metaphor. **Third**, it grounds the framework in the triadic intelligence ontology developed in **Meira** et al. (2026), extending it to the design of educational sequences and assessment regimes. **Fourth**, it offers eight vignettes that translate the framework into concrete daily life across disciplines, modalities and contexts. **Fifth**, it proposes a four-horizon implementation architecture and a seven-level maturity model so that leaders have a navigable pathway. **Sixth**, it maps the political economy of change inside universities, the differentiated trajectories of the Global South, and the risks that justify refusal as a governance capacity in a register that **Meira** (2025b) identifies as defining the present strategic moment.

The paper does not claim a forecast. It does not present a systematic review or original empirical dataset. It is a critical policy synthesis written to be useful to those who must make decisions before all the evidence is in. Where the empirical literature is robust, claims are stated as robust. Where the theoretical case is plausible but underdetermined empirically, claims are stated as plausible propositions. Where the argument is normative, it is identified as a design proposal and accompanied by an account of its refusal conditions. The boundary between description and prescription is policed deliberately, because policy texts often blur it to their detriment.

The intended audience is plural. University leaders — rectors, presidents, provosts, deans — need a navigable account of where to begin without breaking the institution. Faculty and professional staff need a framework that respects their disciplinary judgement and labour. Students and learner advocates need an account of what learner rights mean under AI conditions. Policy-makers and accreditors need a vocabulary that does not reduce educational quality to dashboards. Public commentators need a serious account of what higher education can and cannot become. None of these audiences is served by a text that addresses only one of them.

1.4 A Note on the Hyperobject Metaphor

The term *hyperobject* is used carefully, and not naively. **Morton** (2013) coined it to describe entities so massively distributed in time and space that no single position can perceive them as wholes — climate change, geopolitics, the biosphere. Morton identified five characteristic properties: *viscosity* (the hyperobject sticks to whatever it touches); *nonlocality* (no local manifestation is the whole); *temporal undulation* (the hyperobject distributes itself over multiple timescales); *phasing* (it appears intermittently in any given observation frame); and *interobjectivity* (it produces effects only through relations with other entities). For Morton, hyperobjects are **ontological**; they are not metaphors but real entities whose scale exceeds the human perceptual frame.

Higher education is not climate change. It is a human institution, plural and contestable, with identifiable actors, jurisdictions and histories. To call it a hyperobject in **Morton**'s (2013) strict ontological sense would be to overreach. The term is borrowed here as a heuristic for **three** properties that higher education does share with hyperobjects, in modified form. It is **distributed** across temporal layers that resist synchronous reform: the medieval, the Humboldtian, the mass, the platformised and the AI-mediated coexist within the same institution. It is **composed** of interacting subsystems — pedagogy, assessment, credentials, data, labour, regulation, finance, ecosystems — whose alignment cannot be assumed and whose reform under one heading routinely fails because another heading does not move. And it produces **effects** that exceed any single agent's field of view: a curricular change made by a department can ripple, through accreditation, the labour market, family expectations and platform partners, in ways that no actor can anticipate.

Where systems theory and institutional analysis suffice, they should be used. The contribution of *hyperobject* is rhetorical and disciplinary: it **forces** policy imagination to resist single-surface reforms. Readers who prefer the language of complex sociotechnical systems, of multi-level perspectives on socio-technical transitions, or of institutional theory in the **North** (1990) tradition, may translate the argument into those vocabularies without loss. The metaphor is a working device, **not** an ontological claim. Its limits will be

acknowledged again in section 6, where the fourteen-dimension framework is presented. The substantive payoff is that any reform proposal can be tested against the metaphor by asking: which dimensions move together, which remain in contradiction, and where does the proposal assume more institutional alignment than the historical record warrants?

The metaphor matters, in the end, because higher education reform **repeatedly fails** through a single mode of failure:

> the reformer treats one visible surface — the curriculum, the platform, the credential, the policy — as if it were the whole object.

The point of describing higher education as a sociotechnical hyperobject is to make that failure mode visible, name it and resist it. It is not to suggest that the institution is opaque, beyond reform or beyond accountability. On the contrary: the hyperobject framing is a discipline for ambition, not an excuse for paralysis.

2. Evidentiary Posture

2.1 Critical Policy Synthesis as Method

The article uses a critical policy synthesis. It draws on **five** bodies of evidence, each interrogated for what it can and cannot show.

The **first** body is macro-level data on the structure of higher education, primarily from **UNESCO** and **OECD** (2026) reporting. **UNESCO**'s 2026 framing of higher education as a public mission facing technological, geopolitical, demographic and funding pressures provides the global frame; and the **OECD** (2026) provides the most current institutional analysis of effective and ineffective uses of generative AI in education.

The **second** body is the rapidly growing labour-market literature on AI exposure. **Acemoglu** and **Restrepo**'s (2019) task-based framework of automation as a process that both displaces and reinstates labour provides the theoretical backbone. Eloundou and colleagues (2024) estimate task exposure to large language models. **Brynjolfsson**, **Li** and **Raymond** (2025) report observed productivity gains in a deployed system. **Brynjolfsson**, **Chandar** and **Chen** (2025) identify early employment signals among junior cognitive workers. **Hampole**, **Papanikolaou**, **Schmidt** and **Seegmiller** (2025) analyse hiring and wage effects of AI exposure. **Autor** and **Thompson** (2025) reframe the debate in terms of how AI compresses routine expertise and raises the premium on tacit, integrative expertise. **Massenkoff** and **McCrory** (2026) combine theoretical exposure with observed usage patterns drawn from a frontier-model platform.

The **third** body is the educational research literature on assessment and judgement. **Boud** and **Falchikov**'s (2006) classical argument that assessment must align with long-term

learning, not short-term certification, is the methodological anchor for the assessment-redesign argument advanced in sections 5 and 7. **Bearman**, **Tai**, **Dawson**, **Boud** and **Ajjawi** (2024) update that argument explicitly for the generative AI moment, identifying evaluative judgement as the developmental target that assessment must protect. **Stephenson** and **Armstrong** (2026), in the 2026 HEPI student survey, provide an empirical counterpart that documents how routinely students already use generative AI for tasks that institutions still treat as authentic individual work.

The **fourth** body is the institutional and economic theory of universities. **Trow**'s (1973) tripartite distinction among elite, mass and universal higher education remains analytically useful precisely because countries occupy different points on that trajectory and because mass systems reproduce stratification internally. **Kerr**'s (1963) account of the multiversity as a precarious composition of incompatible communities anticipates many of the contemporary tensions. **Clark**'s (1983) comparative analysis of academic organisation across national systems maps the structural variables — the academic profession, the academic enterprise, system authority — through which AI transformation will be mediated. **Barnett**'s (2011) philosophical defence of *being a university*, **Marginson**'s (2016) recovery of higher education as common good, and **Slaughter** and **Rhoades**'s (2004) analysis of academic capitalism each describe a different layer of the institutional present.

The **fifth** body is policy and industry signal: **UNESCO** and **OECD** (2026) reports; the **Council of the European Union**'s (2022) Recommendation on micro-credentials; the **European Parliament and Council**'s (2024) AI Act; the **NIST** (2023, 2024) AI Risk Management Framework and its Generative AI Profile; Deloitte's 2026 Higher Education Trends (**Clark** et al., 2026); the **Boston Consulting Group**'s analysis (2026); and **UPCEA**'s Predictions 2026 (**UPCEA**, 2025). These are signals, not ground truth; the corporate literature in particular is uneven in methodological rigour and frequently produced inside the AI value chain. The synthesis treats them accordingly: as evidence of where institutional attention is moving and as input to the design proposals that follow.

2.2 Three Levels of Claim

The paper distinguishes three levels of claim, and asks the reader to weigh them differently.

Robust claims are supported by multiple independent sources or by established theory. The claim that AI task exposure is concentrated in cognitive work and varies by task structure is robust across **Eloundou** et al. (2024), **Hampole** et al. (2025), **Autor** and **Thompson** (2025) and **Massenkoff** and **McCrory** (2026). The claim that productivity gains accrue more strongly to less-experienced workers in deployed AI systems is robust within the **Brynjolfsson**, **Li** and **Raymond** (2025) results and consistent with the **Brynjolfsson**, **Chandar** and **Chen** (2025) analysis of junior cognitive labour. The claim that traditional

assessment regimes are increasingly mis-aligned with the conditions under which students actually work is robust across **Bearman** et al. (2024), **Boud** and **Falchikov** (2006) in updated form, and the HEPI survey.

Plausible propositions are theoretically grounded but require comparative empirical research. The claim that the apprenticeship architecture of professional fields is weakening as routine entry-level tasks are absorbed by AI is theoretically well-motivated by the Acemoglu–Restrepo framework, by the Autor–Thompson analysis of expertise, and by the institutional accounts in **Slaughter** and **Rhoades** (2004) and in **Clark** (1983), but it has not been verified across a sufficient comparative sample to be treated as robust. The same applies to the claim that AI-native business models will redistribute scarce human attention toward mentoring, supervision and assessment: the logic is clear, the institutional uptake is uncertain.

Design proposals are normative and strategic. They specify what institutions and policy systems should attempt, with arguments for their plausibility and explicit refusal conditions, but they are not predictions. The nine features of AI-native higher education in section 5, the eight vignettes in section 7, the four implementation horizons in section 9, the ten policy principles in section 10 and the five 2035 scenarios in section 14 are design proposals. They are intended to be argued with, not adopted uncritically. The reader who disagrees with a specific proposal but accepts the underlying framework is the reader the paper most hopes to address.

2.3 Limits and Biases of the Evidence

The strongest empirical evidence on AI and labour markets remains uneven, concentrated in high-income contexts and dominated by United States data. The Eloundou, Brynjolfsson, Hampole, Autor and Massenkoff studies cited above are anchored in US labour-market structures and US-based deployment data. The HEPI survey is United Kingdom-specific. The **OECD** (2026) analysis is broader but is itself filtered through the institutional position of the OECD secretariat. The Global South — where most of the world's students will be educated over the period 2026–2035 — is conspicuously underrepresented in this evidence base, a gap that section 12 addresses directly rather than treating as an afterthought.

The corporate literature carries the additional bias of value-chain interest. Deloitte, the Boston Consulting Group, the Higher Education Policy Institute and other industry sources offer valuable signals but were not produced under disinterested conditions. The paper treats them as input to the description of where institutional attention is moving rather than as evidence for what institutions should do. The **NIST** (2023, 2024) AI Risk Management Framework and the European AI Act are policy artefacts whose implementation is still

being negotiated; they are cited for their architectural intent rather than for proven institutional effects.

The 2026–2035 horizon is treated as a policy horizon, not a forecast window. Forecasting AI capabilities at a ten-year horizon is unreliable on any reading of the recent history of the field; what is more useful, and what policy actors actually need, is a clear specification of the decision space — the choices available now, the risks they create, the values they threaten or advance, the capacities they require and the institutional commitments they imply. That specification is the work of the rest of the paper.

3. The Scale of the Problem: Expansion, Inequality and Legitimacy

3.1 A Planetary System That Has Not Equalised

Higher education is a planetary system. **UNESCO** (2026) reported that global enrolment reached approximately 269 million students in 2024, more than double the level of 2000, with an average gross enrolment ratio of 43 percent and nearly seven million internationally mobile students (**UNESCO**, 2026). The expansion is one of the most consequential institutional facts of the past half century. It has produced new opportunity at scale, new occupational pathways, new research capacity and new public expectations. It has also produced new forms of stratification, new political stakes and new fiscal pressures.

Trow's (1973) framework anticipated much of this. He distinguished **elite** higher education, which serves under 15 percent of the cohort and is organised around the formation of a leadership class; **mass** higher education, which serves 15 to 50 percent of the cohort and is organised around occupational preparation and social mobility; and **universal** higher education, which serves over 50 percent and is organised around continuous adaptation of the population to economic and civic change. Countries occupy different points along this trajectory, and the institutional logics required at each stage are different. A system at universal participation cannot be governed by the assumptions of an elite system, but its institutions are often inhabited by actors trained under the elite paradigm — a source of repeated mismatch that AI will sharpen, not soften.

Expansion has not produced equalisation. Aggregate enrolment growth disguises disparities of quality, completion, employability and finance both between and within national systems. **Slaughter** and **Rhoades**'s (2004) account of academic capitalism describes how massification, under conditions of public funding constraint, generates internal markets that stratify institutions, programmes and academic labour. **Marginson**'s (2016) recovery of higher education as a common good is in part an argument against the

silent privatisation of that good through fee structures and platformised delivery. **Kerr**'s (1963) description of the multiversity as a precarious truce among incompatible communities reads, in 2026, less as historical anthropology and more as a description of what AI is about to stress further.

3.2 Demographic, Fiscal and Legitimacy Pressures

Several mature systems face demographic and fiscal stress. Parts of North America, East Asia and Europe confront shrinking cohorts of traditional-age students. Public funding competes against health, defence, climate, infrastructure and social protection demands. Families question the return on tuition. Employers question whether degrees signal what they once did. UPCEA and Deloitte (**Clark** et al., 2026) report the same signal from inside the market: the centre of gravity of demand is shifting toward shorter, modular, adult and employer-aligned formation. The Higher Education Policy Institute's 2026 student survey (**Stephenson** and **Armstrong**, 2026) shows AI use as routine practice, with implications for what students expect from an institution and what they will pay for.

The pressures are not symmetrical. Selective and prestigious institutions retain pricing power and demand resilience because they sell something AI cannot generate: a positional good underwritten by selectivity, alumni networks, research reputation and accumulated capital. Non-selective and regional institutions, particularly in mature systems, face the sharp end of the demographic and fiscal pressure simultaneously. The Global South faces a different problem: high demand, constrained supply, severe infrastructure and labour limits, and a credentialing system whose international portability is mediated by Northern reference points. The AI moment will not collapse this differentiated geography; it will reshape the boundary conditions inside each segment.

Legitimacy is the deeper stake. The **Boston Consulting Group**'s analysis (2026) frames the AI moment as a once-in-a-generation opportunity for higher education, but its argument depends on institutional credibility being preserved through the transition. If universities are perceived to be **either** captured by platform providers, defensive against change **or** unable to certify what they claim to certify, that credibility will erode quickly. The legitimacy problem is therefore not that higher education is failing everywhere; it is that the default bundle of higher education is becoming less coherent as a single package at the moment when its legitimacy is most exposed.

3.3 The Unbundling of Higher Education

Universities have historically sold a **bundle**: instruction, certification, formation, social capital, employability signalling and research access. Each of these components is now under pressure from independent forces — open content, alternative credentials,

employer-led training, AI tutoring, contract research, online platforms — and AI intensifies the unbundling by changing the cost and availability of the activities inside each component. The unbundling is **not** the consequence of any single decision; it is the aggregate **effect** of *many actors optimising at component level while the institution continues to depend on the integrity of the bundle*.

The platform analysis in **Meira** (2025a) identifies a recurrent dynamic in such moments. When a general-purpose digital infrastructure lowers the cost of orchestration across previously bundled functions, value migrates from the original bundle to specialised providers and to whichever actor controls the orchestration layer. The same dynamic that has reshaped retail, finance, media and mobility now reaches the bundle of higher education. Whether the orchestration layer is captured by a small number of global platforms, distributed among public actors, or shared in some hybrid arrangement is the political question of the next decade.

The **defensible** value of higher education after generative AI lies **not in** content delivery, which AI undermines, **but in** the formation and credible certification of capabilities that require judgement, accountability, social learning, sustained inquiry and verified human performance.

> This reframes employability **not as** a list of skills **but as** the architecture by which novices become trustworthy professionals — an architecture that **Meira**, **Neves** and **Braga** (2025) analyse as the transition from programmed labour to meta-cognitive orchestration.

The Drucker-era (**Drucker**, 1959) image of the knowledge worker, formalised in the late 1950s and consolidated through the second half of the twentieth century, was already an answer to the previous unbundling: the migration of value from physical to cognitive operations. The present moment is an unbundling of the cognitive bundle itself, and the institutional response cannot be a faster version of what universities did between 1959 and 2022.

4. AI, Labour Markets and the Weakening Graduate Entry Contract

4.1 From Task Exposure to Observed Effects

Early studies of large language models estimated exposure by comparing occupational tasks with model capabilities. **Eloundou** and colleagues (2024) argued that a large share of the United States workforce could have at least **ten** percent of work tasks affected by LLMs, and that around **nineteen** percent of workers could have half of their tasks exposed.

Exposure was concentrated in higher-education and higher-income occupations — the inversion of the classical pattern in which automation hit routine manual work. The result was provocative on first reading and remains the starting point for the policy literature, but it is best read as a feasibility map rather than an employment forecast.

Subsequent task-level work moved from feasibility to observed effect. **Brynjolfsson**, **Li** and **Raymond** (2025) reported productivity gains concentrated among less-experienced workers in a customer-service deployment, with implications for the apprenticeship structure of professional careers that section 4.2 develops. **Brynjolfsson**, **Chandar** and **Chen** (2025) identify early employment signals among junior cognitive workers — the "canaries in the coal mine" — that suggest compositional effects within occupations rather than wholesale displacement. **Hampole**, **Papanikolaou**, **Schmidt** and **Seegmiller** (2025) analyse hiring and wage effects of AI exposure, finding patterns consistent with task rather than worker substitution.

Acemoglu and **Restrepo**'s (2019) broader framework of automation as a process that displaces and reinstates labour is the theoretical backbone for reading these results. AI is, in their terms, a force of compositional change rather than a single substitution shock; the policy question is which tasks are displaced, which are reinstated in new forms, and which new tasks are generated by the technology itself. **Autor** and **Thompson** (2025) formalise a useful corollary: AI is more likely to compress the value of routine expertise while raising the premium on tacit, contextual and integrative expertise. The implication for higher education is direct. If AI compresses routine expertise, the curriculum that produced that expertise must be redesigned; if AI raises the premium on tacit, integrative judgement, the formation of that judgement must be the institution's defensible product.

Massenkoff and **McCrory** (2026) combine theoretical feasibility with real-world usage patterns from a frontier-model platform. They find limited evidence of increased unemployment in highly exposed occupations, but compositional and entry-level effects that warrant policy attention. Taken together, the labour-market evidence suggests neither apocalypse nor business as usual: a redistribution of where cognitive value is produced, with the most disruptive effects concentrated at the entry of professional careers. The **OECD**'s (2026) study draws a similar conclusion in the educational register: the central problem is not unemployment of graduates but the **silent erosion** of the developmental pathway through which graduates became experts.

4.2 The Apprenticeship Paradox

For higher education, the central issue is not a mass-unemployment forecast. It is the graduate entry contract. Universities have historically certified the readiness of young adults **for** entry-level professional roles.

But many entry-level roles consisted precisely of the routinised cognitive tasks that AI now performs adequately: document review, code scaffolding, summary preparation, draft analysis, customer triage, paralegal research and elementary data work. Junior lawyers learned by reviewing documents; junior engineers learned by writing code; junior consultants learned by building models; junior analysts learned by cleaning data.

These tasks were not just labour. They were apprenticeships.

Autor and **Thompson**'s (2025) expertise framework helps explain why this matters: automation does not simply reduce the quantity of tasks in an occupation; it can alter the expertise required by the remaining task bundle, sometimes removing lower-expertise support work and leaving behind more expert problem-solving and decision-making tasks.

This is the apprenticeship paradox.

Universities are asked to prepare students for high-level judgement, but workplaces used lower-level tasks as the mechanism through which judgement was learned.

If those mechanisms vanish, judgement must be developed inside the educational architecture itself, or alongside it through deliberate work-integrated learning.

The paradox is not entirely new. **Polanyi**'s (1966) analysis of the tacit dimension of expert knowledge established that explicit instruction can never substitute for the apprenticeship-like accumulation of pattern recognition in practice.

What is new is the **loss** of the institutional infrastructure that delivered the apprenticeship as a by-product of paid work.

The implication is that employability cannot be reduced to a list of skills printed on a transcript. It must be the documented, verifiable capacity to do meaningful professional work — including the work of supervising AI — under domain-specific **conditions**.

The argument is consistent with **Meira** et al. (2026), which formalises meta-cognitive orchestration as the defining professional competence of the twenty-first century and locates educational design as its principal site of formation. The **triadic framework** — individual intelligence (cognition, emotion, creativity), social intelligence (networks, institutions, collectives) and artificial intelligence (functional agency without phenomenology) — provides the analytical scaffolding for designing what now has to happen inside the curriculum because it can no longer be assumed to happen on the job.

The institutional response **is not** to imitate the workplace.

> Universities **cannot** replace, at scale, the apprenticeship that the firm once provided.
>
> They **can**, however, redesign their architecture so that *authentic, supervised, integrative practice is built into the curriculum*: clinics, studios, capstones, simulations, fieldwork, internships, project-based learning under expert oversight, with explicit attention to the supervision of AI.

This is **expensive** in human time, which is precisely why the architecture of AI-native higher education — the subject of sections 5, 7 and 9 — *turns on the reallocation, not the reduction, of academic labour*.

4.3 The Mirror Problem Inside the University

The same paradox plays out **inside** the university. Academic expertise has long been reproduced through junior and often invisible work: grading drafts, preparing lectures, tutoring small groups, conducting literature reviews, maintaining laboratories, advising students, supporting admissions and administering committees.

> Much of this work is performed by graduate students, adjuncts, postdoctoral fellows and professional staff under conditions that **Slaughter** and **Rhoades** (2004), in their analysis of **academic capitalism**, document as **increasingly precarious**.
>
> AI now reduces the marginal cost of much of this work.
>
> The risk is not only that adjuncts and graduate students are displaced; it is that the institutional pathways through which mid-career academics learned to teach, supervise and govern are **silently eroded**.

A university that asks its students to become meta-cognitive orchestrators while organising its own labour as programmed compliance is in a position of performative **contradiction**.

> The contradiction is not abstract. It plays out in the workload models that pay faculty for measurable activities (lectures delivered, papers published, students supervised) while leaving the orchestration work — the design of assessment, the supervision of teaching assistants, the maintenance of the learning architecture — unrecognised. It plays out in the promotion criteria that reward research output while treating teaching reform as service. It plays out in the procurement decisions that buy platforms which deskill rather than augment the educator. **Meira**, **Neves** and **Braga** (2025) name this risk explicitly under the heading of **digital Taylorism**: the use of AI to fragment academic judgement into dashboards and prompts, with educators reduced from designers of learning to monitors of platform outputs.

The countermove is **reflexive transformation**. The institution must apply to its own labour the framework it asks its graduates to internalise.

> Faculty workload models must recognise design, supervision, assessment and orchestration.
>
> Promotion criteria must reward the redesign of teaching and assessment alongside research.
>
> Procurement must be evaluated for whether it augments or deskills academic work.
>
> Professional staff must be recognised as carriers of institutional knowledge rather than as compliance functions.

The reflexive condition is not a moral demand but an **operational one**: a university that fails it cannot credibly certify what it claims to teach.

4.4 The Equity Stakes of the Entry Contract

The graduate entry contract has equity implications that are easily underestimated when AI policy is written in elite institutions. Students from elite institutions retain access to internships, alumni networks, research experiences and high-status human mentorship. Students in less-resourced institutions may receive lower-quality AI tools, weaker human support and credentials that no longer signal what they once did. The **reproduction** of advantage through informal networks of opportunity is well documented in the sociology of higher education from Bourdieu onwards; AI provides a **new** mechanism through which the same reproduction can occur under the rhetoric of access. Without explicit policy intervention, AI will **widen** the gap between the prestige bundle and the access bundle.

The risk has a specific institutional shape. Elite institutions may deploy AI as a **complement** to intensive human mentorship: the student gets the AI tutor and the seminar with the senior professor. Mass-access institutions, under fiscal pressure, may deploy AI as a **substitute** for human support: the student gets the AI tutor and not much else. **Stephenson** and **Armstrong** (2026) document, in the 2026 HEPI survey, exactly this differential expectation: students at well-resourced institutions report AI as a supplement to teaching; students at under-resourced institutions report AI as a replacement for teaching they cannot otherwise access. The mass-access pattern is rationalised as inclusion but functions as a **downgrading** of the formation those students receive.

A policy framework for AI-native higher education must therefore make equity a **design constraint**, not a downstream concern. The principle is operationalised in section 10 through specific instruments — equity-disaggregated indicators, public procurement requirements, public infrastructure provision, refusal conditions on uses of AI that

compound stratification — and in section 12 through a detailed treatment of the Global South as a site of theory rather than a site of adoption. The point is that equity is not an additional thing that AI-native higher education must do; it is one of the principal tests by which the AI-native claim is judged.

5. From AI-Enabled to AI-Native Higher Education

5.1 Conceptual Discipline Against Marketing Capture

Conceptual discipline is necessary because the language of AI adoption is easily **captured by marketing**. Many institutions will describe themselves as AI-ready, AI-enhanced, AI-powered or AI-enabled. These labels may reflect real work, but they are insufficient and often misleading.

> An **AI-enabled** university **adds tools** to an inherited architecture.
>
> An **AI-native** university **redesigns** the architecture around the existence of ubiquitous AI as cognitive infrastructure. Its question **is not** which tools to license, **but what** learning, assessment, credentials, data, academic work, governance and business models look like when AI is part of the substrate of cognition.

The distinction has historical precedent. When electricity was a new general-purpose technology, factories **first** electrified their existing steam-powered architecture by replacing the central steam engine with a central electric motor; productivity gains were **marginal**. The substantive transformation occurred decades later, when factory layouts were **redesigned** around distributed motors at the point of use.

> **The educational analogue is direct**. The AI-enabled institution replaces the central steam engine; the AI-native institution redesigns the factory. The historical lag in the electrification case — measured in decades — is also a warning. Institutions that wait for clarity before redesigning **will be redesigned** by their environment.

The platform analysis in **Meira** (2025a) identifies the same dynamic at the level of value (network) chains. When a general-purpose infrastructure **lowers** coordination costs across previously bundled functions, the institutional question is whether the incumbent will be the **orchestrator** of the new architecture or the **supplier** inside someone else's.

> Universities that remain AI-enabled — adding tools to inherited architecture — are positioning themselves as **suppliers** inside **platforms** designed by others. Universities that become AI-native are positioning themselves to be the **orchestrators** of the educational layer of their own institutional **ecology**.

5.2 Nine Features of an AI-Native Institution

First, AI literacy and meta-cognitive orchestration are **general** educational requirements, not specialist electives.

> Every graduate understands what AI systems can and cannot do, how they fail, how they bias decisions, how they should be supervised, and how to compose individual, social and artificial intelligences in concrete tasks. The thirteen foundational principles re-inventoried in **Meira** et al. (2026) provide a curricular template for what this literacy must contain at minimum: **from** the irreducibility of phenomenological human experience, **through** the role of social validation in epistemic claims, **to** the design principles for human-AI collaboration under high-stakes conditions.

Second, assessment is redesigned around authentic performance, supervised evidence and process visibility.

> The unsupervised written assignment cannot remain the **dominant** evidence of **individual** competence. The argument is developed by Bearman and colleagues (2024) in explicit terms, but it rests on the older principle articulated by **Boud** and **Falchikov** (2006): assessment must align with long-term learning, not the convenience of short-term certification. Oral defences, live performances, simulations, portfolios, decision logs and process traces become as central as the artefact itself. The **shift** is not a methodological detail; it is the substantive content of what assessment now has to do.

Third, curriculum is layered and modular.

> Foundations of reading, writing, reasoning and method are **protected and re-emphasised**. AI and data fluency are **distributed** across disciplines. Domain expertise is **integrated** with situated judgement. Modules are stackable, portable and assembled into coherent formation. The **Council of the European Union**'s (2022) Recommendation on micro-credentials provides an institutional template for the modularity that does not collapse into fragmentation. Modularity is a **means**, not an end; the end is **coherent** formation under conditions of differentiated demand.

Fourth, pedagogy is composed, not replaced.

> A student may interact daily with an AI tutor, weekly with peers, periodically with mentors and intensively with faculty during studios, laboratories, oral defences or projects. The educator's role **shifts** from delivery to orchestration of human and machine sequences. The composition logic is the pedagogical face of the meta-cognitive orchestration framework developed in **Meira**, **Neves** and **Braga** (2025):

what students learn is what their educators are simultaneously learning to do at professional scale.

Fifth, learner data governance becomes an academic function rather than a compliance afterthought.

Learning records, AI interaction traces, assessment evidence and credential data are part of educational infrastructure and therefore **require** rights, oversight, audit and portability. The European AI Act (**European Parliament and Council**, 2024) treats several education and vocational uses as high-risk and imposes documentation, transparency and human-oversight obligations accordingly. The **NIST** (2023, 2024) AI Risk Management Framework and its Generative AI Profile (**NIST**, 2023, 2024) provide a vocabulary that institutions can adopt for procurement and governance **without** locking themselves to specific vendors.

Sixth, credentials are unbundled and recomposed.

Degrees, modular credentials, micro-credentials, badges and verified records of practice are connected by coherent learning frameworks, quality assurance and learner protection. The micro-credentials recommendation cited above is one template; the structural argument is that credentials should travel **with** the learner, certify what the learner can demonstrably **do**, and stack into coherent pathways **without** depending on the integrity of a single institutional gatekeeper. The political economy of credential recognition is examined in section 11.

Seventh, business models are reorganised around scarce human attention.

AI **reduces** the marginal cost of routine instruction and administration; the institution **invests** the released time into mentoring, feedback, supervision and research formation, not into cost reduction alone. The strategic risk identified by **Slaughter** and **Rhoades** (2004) — that academic capitalism converts every efficiency into a market signal — applies with new force under AI conditions. The **countermove** is to commit, *ex ante*, that *the gains from automation will be reinvested in human attention*. Without such a commitment, AI-native rhetoric will mask AI-driven adjunctification.

Eighth, the ecosystem is treated as design space.

Employers, professional bodies, governments, civic organisations and other educators participate in **co-designing** pathways, assessment and credentials under shared standards rather than as customers of an opaque product. This implies a different posture toward the platform economy. The orchestration logic of **Meira** (2025a) is reapplied at the institutional layer: universities **become** orchestrators of

educational ecosystems rather than passive recipients of platform terms set elsewhere.

Ninth, the university applies the framework to itself.

Faculty, librarians, learning designers, advisors, technologists, assessment specialists, accessibility experts and administrators **redesign their own work** as meta-cognitive orchestration. AI-native transformation that exempts the institution's own labour from change **is not** AI-native. *The reflexive condition is the test of seriousness*. It is what distinguishes an AI-native claim from an AI-themed prospectus.

5.3 The Triadic Foundation: Individual, Social and Artificial Intelligence

The nine features above presuppose a conceptual scaffold that **Meira** et al. (2026) develops in detail. The triadic framework identifies three **irreducible** kinds of intelligence whose composition is the substantive content of meta-cognitive work.

Individual intelligence encompasses cognition, emotion and creativity: the embodied, phenomenologically grounded capacity of a human person to perceive, judge, feel and create. It is **not** reducible to performance on benchmarks. Its formation is the historical purpose of the university, and its **protection** is the educational **imperative** under AI conditions.

Social intelligence is the property of networks, institutions and collectives: the distributed capacity to coordinate, validate, deliberate and act through structures that **no** individual contains. It is what makes scientific knowledge **possible**, what makes professional standards **enforceable**, and what makes democratic governance **more** than aggregated preference.

Artificial intelligence, in the precise sense relevant here, is the **functional agency** exhibited by *computational systems that perform operations historically associated with cognition without the phenomenological substrate* that grounds individual intelligence. The semantic distinction is consequential. To say that AI is a third kind of intelligence is **not** to ascribe to it experience, understanding or moral status.

It is to **recognise** that, functionally, AI systems now occupy a cognitive role that **cannot** be reduced to either individual or social intelligence and whose composition with the other two is the **new** design problem. The framework draws on, but does not retread, the question **Turing** (1950) posed: *it sets aside whether machines can think and asks how their functional agency composes with the other two intelligences in tasks of consequence*.

The composition of the three intelligences produces a synergistic space that **Meira** et al. (2026) formalises as a **three-dimensional lattice** with characteristic regions:

> tasks where individual intelligence dominates, tasks where social intelligence dominates, tasks where artificial intelligence dominates and — most consequentially — regions of **super-additivity** where the three together accomplish what none alone could.

The pedagogical implication is that an AI-native curriculum **must** teach students to operate competently in each region and to navigate among them. Recognising which task lives in which region — which decision requires phenomenological judgement, which requires social validation, which can be delegated to functional agency under supervision — is a **meta-cognitive** competence, not a domain skill.

The framework also clarifies what **cannot** be delegated. Phenomenological experience cannot be delegated; social legitimacy cannot be delegated; ultimate accountability for decisions taken in the world cannot be delegated. These constraints are not philosophical asides. They define the **boundary** conditions for the design of educational sequences, assessment regimes and professional roles in the AI-native institution. The vignettes in section 7 are deliberately constructed to make those boundaries visible at the level of daily institutional life.

5.4 A Curriculum in Five Layers

The AI-native curriculum can be described in five layers.

> The **first** is foundational intellectual capability: reading, writing, mathematical reasoning, historical consciousness, scientific method, argumentation, interpretation and communication. These foundations become more important, not less, because they are what makes AI outputs evaluable.
>
> The classical educational tradition — **Newman**'s (1852) *idea of a university* as a place of universal knowledge, **Barnett**'s (2011) philosophical defence of the university as a site of critical interpretation, the long German and French traditions of *Bildung* and *culture générale* — is **not** a quaint historical residue. It is the **precondition** of competent AI use. A student who cannot read carefully, write precisely or reason rigorously **cannot** evaluate what an AI produces, and is therefore at the mercy of it.
>
> The **second** layer is AI and data fluency: an understanding of model behaviour, training data, evaluation, prompt and tool use, hallucination, bias, governance and refusal conditions. The fluency is **not** technical specialism; it is the operational and critical algorthimic literacy of a citizen and professional who will spend the rest of

life working alongside such systems. Eloundou and colleagues' (2024) exposure analysis indicates that this fluency is now relevant across virtually every cognitive occupation. The **Bearman** et al. (2024) analysis of evaluative judgement specifies what that fluency must enable at the level of professional practice: the **capacity** to evaluate, justify and defend the use of AI outputs in a defined domain.

The **third** layer is triadic orchestration: the practical capacity to compose individual, social and artificial intelligences in concrete tasks, *with attention to which intelligence is responsible for what*. The capacity is the educational counterpart of the meta-cognitive orchestration analysed in **Meira**, **Neves** and **Braga** (2025) at the level of professional work. Its pedagogical formation requires **specific** design: tasks that explicitly require the student to decide (whenever possible) which intelligence to deploy where, with what supervision, under what accountability, and to defend that allocation under examination. The decision log in the business school capstone of section 7.3 is a concrete instance.

The **fourth** layer is domain depth: substantive expertise in disciplines or professional fields, **recomposed** in light of which operations AI changes. Domain expertise does not disappear in the AI-native curriculum; it becomes **more specific**. The **Autor** and **Thompson** (2025) analysis of expertise compression and elevation indicates which parts of domain expertise are at risk and which are reinforced. The **institutional** task is to redesign disciplinary curricula so that the AI-resistant frontier — tacit, integrative, contextual judgement — is the **substantive** content of the domain experience.

The fifth layer is integrative and civic formation: capstones, studios, clinics, public projects, research apprenticeships and reflective practice that build the judgement and accountability AI alone **cannot** certify. The layer is what makes the previous four layers cohere as **formation** rather than as a portfolio of skills. **Marginson**'s (2016) common-good framing of higher education and **Kerr**'s (1963) multiversity analysis converge here: *the university's contribution to public reason is realised in the **integrative experience**, not in any one layer of preparation*.

5.5 Pedagogy as Composition

Pedagogy in an AI-native institution is a composition problem rather than a delivery problem. A course is best understood as a sequence: preparation supported by AI; practice with formative feedback; collaboration in peer groups; mentoring at human pace; live, supervised performance; reflection that documents what was learned, what was delegated and what was verified. The educator **designs** the sequence and **orchestrates** it; AI participates as a **peer** tool with explicit constraints; assessment **demonstrates** what the

student can do in conditions the institution can defend. The sequence is generic; its instantiation in a particular course depends on the domain, the level, the cohort and the institutional context.

Three design heuristics structure the composition.

> **First**, AI is most educationally valuable when it is **embedded** in a sequence rather than substituted for a step in one. A tutor that produces the answer for the student short-circuits the developmental purpose; a tutor that asks the student to predict, justify and compare with model answers develops it.
>
> **Second**, human time is the **binding** resource and should be **invested** where it does the most developmental work: in mentoring, supervision, integrative discussion, oral defence and authentic feedback. **Bearman** et al. (2024) argue for this allocation as the institutional consequence of taking evaluative judgement seriously.
>
> **Third**, assessment must be designed **before** instructional design, not after, because what is assessed is what is learned; the alignment principle in **Boud** and **Falchikov**'s (2006) analysis is structural, not pedagogical, and applies with new force when unsupervised written work is no longer a reliable measure.

The composition logic implies a different definition of what counts as **good** teaching. Good teaching in an AI-native institution **is not** the live performance of explanation, which AI does at low marginal cost. It is **the design** of a sequence in which the student becomes capable of performing the explanation themselves. The educator's craft is the craft of orchestration. The **orchestration** includes the human teaching assistants, the peer learners, the AI tutor, the simulation environment, the reflective journal and the oral examination — all of which the educator must compose with **pedagogical intent**. The implication for academic labour is examined in section 4.3 above and operationalised in section 9.

5.6 What AI-Native Is Not

Conceptual clarity demands a list of **negations**.

> AI-native higher education *is not the replacement of teachers by chatbots*. The replacement narrative is empirically unsupported, pedagogically incoherent and politically dangerous. It is not the reduction of universities to training platforms. The training platform is a legitimate institutional form but not the only one, and not the form that can carry the civic and scientific functions universities also perform. It is not surveillance-based learning analytics. **Selwyn**'s (2019) critical analysis of robotic-teacher rhetoric is one strand of the broader argument that surveillance dressed as personalisation degrades both learning and rights.

It is not a lowering of standards. Authentic assessment, supervised evidence and process-visible work are more demanding than the unsupervised written assignment they replace, not less. *It is not the privatisation of public mission*. The platformisation of educational infrastructure under private terms is the principal political-economy risk of the transition, and section 13 specifies refusal conditions accordingly. *It is not the substitution of human relationships by automated coaching*. The relational substrate of formation cannot be automated, and any institutional design that attempts to do so commits the categorical mistake that **Meira** et al. (2026) identifies as *conflating functional agency with phenomenological intelligence*.

It is not "efficiency" as a euphemism for cost-cutting. The AI-native institution may be more productive, but its **principal commitment** is the reinvestment of released time into the human attention that defines its value.

Each negation is an institutional commitment to be defended explicitly in governance, procurement, assessment and labour. The negations **are n**ot rhetorical flourishes. They are the boundary conditions of the design proposal; without them, AI-native rhetoric collapses into the AI-themed adjunctification that section 4.3 identifies as the principal internal risk.

6. Higher Education as a Sociotechnical Hyperobject

6.1 The Reform Failure Mode the Metaphor Names

Higher education reform repeatedly **disappoints** because reformers often treat one visible surface as the whole object. A university may introduce digital platforms **without** changing pedagogy. A regulator may require learning outcomes **without** changing funding or assessment. An employer coalition may launch alternative credentials **without** changing accreditation. A faculty group may redesign assessment **without** changing workload models. Each reform may succeed locally and fail systemically because the surrounding architecture does not move. The pattern is sufficiently consistent across national systems and historical episodes that it deserves a name. The hyperobject framing offers one.

Morton's (2013) original concept is ontological and ecological. Hyperobjects are real entities whose scale exceeds the human perceptual frame: they are viscous (they cling to whatever they touch), nonlocal (no local manifestation is the whole), temporally undulating (they distribute themselves across multiple timescales), phased (they appear intermittently in any given observation frame) and interobjective (they produce effects only through relations with other entities). The properties were developed in dialogue with planetary phenomena: climate change, the biosphere. Their migration into discussions of human

institutions deserves **caution**; the migration must be a controlled borrowing, not an undisciplined analogy.

Higher education is **not** climate change. It is a human institution, plural and contestable, with identifiable actors, jurisdictions and histories. To call it a hyperobject in **Morton**'s (2013) strict ontological sense would be to overreach. The term is borrowed here as a heuristic for three properties that higher education does share with hyperobjects, in modified form.

> It is **distributed across temporal layers** that resist synchronous reform: the medieval, the Humboldtian, the Newmanian, the multiversity, the platformised and the AI-mediated coexist within the same institution. **Clark**'s (1983) comparative analysis of academic organisation across national systems documents this temporal layering in structural terms.
>
> It is composed of **interacting subsystems** whose alignment cannot be assumed: the fourteen dimensions identified in section 6.2 are the most explicit articulation.
>
> It produces effects that **exceed any single agent's field of view**: a curricular change made by a department can ripple, through accreditation, the labour market, family expectations and platform partners, in ways that no actor can anticipate.

Where systems theory and institutional analysis suffice, they should be used. **North**'s (1990) distinction between institutions (rules of the game) and organisations (the players) is the analytical backbone of section 8. **Slaughter** and **Rhoades**'s (2004) academic capitalism describes the political-economy layer. **Marginson**'s (2016) common-good framing describes the normative layer. **Kerr**'s (1963) multiversity describes the historical-structural layer.

These vocabularies are necessary, and the hyperobject framing does not replace them. Its contribution is rhetorical and disciplinary: it forces policy imagination to resist single-surface reforms by holding the distributed, multi-timescale, alignment-requiring character of the object in view. Readers who prefer to translate the argument back into the language of complex sociotechnical systems lose nothing analytically; they may lose only the discipline of ambition that the borrowed term imposes.

6.2 Fourteen Aligned Dimensions

Higher education as a hyperobject can be analysed along fourteen aligned dimensions, each carrying its own policy question and its own AI-native design demand. The list is not a taxonomy. It is a checklist of alignment requirements. A reform proposal can be tested by asking, for each dimension, which way it pushes and which it leaves untouched. A micro-credential strategy that ignores assessment, academic work and accreditation will fail. A

campus AI policy that does not touch business models, ecosystems and academic-work redesign will not endure. The discipline of the framework is to **require**, for every reform proposal, an explicit answer to the question: which dimensions are aligned, and which are left in contradiction?

The fourteen dimensions are: purpose; epistemic order; curriculum; pedagogy; assessment; technology; space–time; economics; institutional form; regulation; academic work; ecosystems; culture and identity; and ethics and public value. Each has been explored in some form in the analytical literature:

1. Purpose by **Newman** (1852), **Barnett** (2011) and **Marginson** (2016)
2. Epistemic order by the philosophy of science
3. Curriculum by the comparative-curriculum literature
4. Pedagogy by educational research
5. Assessment by **Boud**, **Bearman** and colleagues (2024)
6. Technology by sociotechnical studies
7. Space–time by the architecture and sociology of educational space
8. Economics by **Slaughter** and **Rhoades** (2004) and the higher education finance literature
9. Institutional form by **Clark** (1983) and the comparative-systems literature
10. Regulation by the quality-assurance literature
11. Academic work by **Meira**, **Neves** and **Braga** (2025) and the academic-labour literature
12. Ecosystems by the innovation-systems literature
13. Culture and identity by the historical and sociological literature on the university
14. Ethics and public value by **Marginson** (2016) and the ethics of education literature

Table 1 presents the policy question and AI-native design demand for each dimension.

Table 1. Fourteen dimensions of higher education as a sociotechnical hyperobject

Dimension	Policy question	AI-native design demand
Purpose	Why should higher education exist under AI conditions?	Rebalance formation, research, social mobility, public reason, professional capability and planetary responsibility.
Epistemic order	What counts as valid knowledge when outputs are machine-generated?	Teach verification, provenance, confidence, uncertainty, domain standards and social validation.
Curriculum	What should all graduates learn?	Integrate foundations, AI and data fluency, phygital and triadic orchestration, domain expertise and situated judgement.
Pedagogy	How should learning occur?	Compose AI-supported practice, reflection, peer learning, human mentoring, projects, laboratories and seminars.
Assessment	How can competence be trusted?	Shift from artefact verification to competence verification through authentic, supervised and explainable evidence.
Technology	Who controls the infrastructure?	Require interoperability, accessibility, cybersecurity, audit rights and algorithmic transparency, open standards and public-interest options.
Space–time	Where and when does learning occur?	Combine campus, workplace, community, home, networked and asynchronous learning without losing belonging.
Economics	Who pays, what costs fall and what remains scarce?	Use AI to reduce routine costs while protecting human-intensive feedback, mentoring and evaluation.
Institutional form	Which organisations can legitimately educate and certify?	Enable diverse providers while protecting quality, equity and public value.
Regulation	How are quality and risk governed?	Modernise accreditation, data protection, AI governance, credential recognition and professional licensure.

Academic work	What becomes of faculty and educational labour?	Redesign roles, contracts, workload, continuous professional development, promotion and academic freedom for meta-cognitive work.
Ecosystems	How does higher education connect to work and society?	Build employer, civic, professional, research and public-sector partnerships around verified capability.
Culture and identity	What rites, communities and symbols matter?	Preserve belonging, intellectual community and institutional memory while redesigning delivery.
Ethics and public value	What should not be automated or commodified?	Define refusal conditions, equity requirements, learner rights, democratic oversight and sustainability.

6.3 The Informaticity Substrate

Underneath the fourteen dimensions sits a substrate whose theoretical structure **Meira** and **Meira** (2026) analyses in formal terms: **informaticity**. The substrate has three axes, each named after the foundational paper that defined it.

> The **computability axis** (**Turing**, 1936) specifies what can be effectively calculated by a finite procedure.
>
> The **communicability axis** (**Shannon**, 1948) specifies what can be reliably transmitted across a noisy channel.
>
> The **controlability axis** (**Wiener**, 1948) specifies what can be steered through feedback.

The convergence of the three axes into a **single** deployable infrastructure — networks, platforms, models, devices — is the historical event that distinguishes the present.

The framework distinguishes informaticity-T, the theoretical possibility space defined by the three axes, from informaticity-P, its historical realisation through platforms, networks and sociotechnical systems. Generative AI is a **moment** in the development of informaticity-P, not a step beyond informaticity-T.

> The distinction matters for policy because it locates the AI moment **inside** a longer arc rather than treating it as a unique exception to all previous experience.
>
> Higher education has lived through earlier moments of the same arc: the printing press, broadcast media, the internet, the smartphone, the platform economy. Each reorganised the cost structure of cognitive operations; each prompted institutional

debate; each was eventually metabolised into the institutional architecture without dissolving it. Generative AI is the **most consequential** of these moments because it touches all three axes simultaneously and at scale.

The phygital framing of **Meira** and **Meira** (2026) adds a further specification. Contemporary informaticity is realised in a triaxial space: the Physical (P), the Networked Digital (D) and the Networked Social (S).

> The **Physical** is ontologically prior — the only dimension capable of independent existence and the structural precondition for the others.
>
> The **Digital** is grounded in the computability axis; the **Social** in the communicability axis; their recursive interdependence in the controlability axis.

The three dimensions are irreducible to one another, but they are not symmetrically related. The asymmetry has institutional implications. Universities are above all Physical institutions: they have *campi*, bodies, laboratories, archives, libraries. The temptation to relocate them entirely to the Digital and Social dimensions **ignores** the ontological priority of the Physical and the formative role of **embodied presence** in the development of individual and social intelligence.

The substrate framing produces a useful diagnostic. **AI-enabled** institutions tend to act on one or two dimensions of the substrate without changing the others: they digitise **without** changing the social architecture, or they reorganise the social architecture **without** addressing the physical and digital substrate.

AI-native institutions **redesign across all three dimensions** of the phygital substrate **and all three axes** of informaticity, with explicit attention to the asymmetric relation among them. The fourteen-dimension framework of section 6.2 is therefore **not** floating in conceptual space; it sits on a specific substrate whose properties have been formally analysed.

6.4 The Red-Swan Register of the Transition

The transition itself has a register that classical strategic planning was not designed to handle. **Meira** (2025b) develops a typology of strategic ruptures that exceed the Black Swan vocabulary inherited from Taleb. Red Swans are **not** low-probability events with high consequences. They are events whose *ontology* the inherited strategic frame cannot accommodate.

> The typology specifies five facets: **hyperconnectivity** (systemic fragility in a coupled world), **synchronicity** (cascading rupture across coupled temporalities), **irreversibility** (lock-in and entropy that prevent return to prior states), the

invisibility paradox (institutionalised blindness to weak signals) and paradigmatic **transcendence** (ontological displacement and epistemic rupture).

Each facet applies to the higher education transition in specific form.

Hyperconnectivity: a single policy change in a major regulatory market — the EU AI Act, the United States executive landscape, China's educational AI policy — propagates through global platforms and partnerships in ways that no single institution can control.

Synchronicity: the demographic cliff, the assessment-credibility crisis, the labour-market reconfiguration, the platform-capture pressure and the political contestation of higher education are arriving on coupled timescales rather than as separable problems.

Irreversibility: once a generation of students has been certified under conditions of unsupervised AI use, the credentials they hold cannot be retrospectively re-certified, and the institutional consequences will be lived for decades.

Invisibility: institutional governance is structurally inclined to underweight signals that originate outside its established decision channels — student practice, platform infrastructure decisions, junior-labour displacement — until the signals are too large to ignore.

Paradigmatic transcendence: the AI transition does not preserve the categories on which inherited educational policy was built (the bounded course, the individual assignment, the place-based campus, the time-defined degree), and reform inside those categories increasingly produces gestures of compliance rather than substantive change.

The **strategic implication** is not that higher education is doomed but that the strategic vocabulary universities have used to plan their reform **must be** supplemented by a vocabulary capable of holding red-swan conditions in view.

Strategic humility, pluralism of options, institutional reflexivity, constructive irrationality (designing for the impossible) and *temporal intelligence beyond forecasting* — the five capacities **Meira** (2025b) identifies — are the cognitive infrastructure of credible governance under the present conditions.

The implementation architecture proposed in section 9 is designed to be compatible with that infrastructure rather than to displace it.

7. Day in the Life: Eight Design Vignettes

Abstract frameworks leave readers tactically adrift. The idea of AI-native higher education becomes tangible only when it is rendered as institutional daily life. The following vignettes are hypothetical. They are not models to be copied. They are design probes that show what AI-native composition can look like across disciplines, modalities and contexts. Each vignette deploys the triadic intelligence framework of **Meira** et al. (2026) and the meta-cognitive orchestration logic of **Meira**, **Neves** and **Braga** (2025) in concrete pedagogical form, and each anchors the assessment design to the evaluative-judgement framework of Bearman and colleagues and the alignment principle of **Boud** and **Falchikov** (2006). The vignettes are written from the participant's perspective — the student, the seminar, the team, the faculty member — because the unit of analysis of any educational reform is, in the end, a person learning.

7.1 The First-Year Biology Student

Mariana is a first-year biology student at a public research university.

> Her week begins with an AI-supported preparation sequence keyed to the next laboratory. The system does not summarise the textbook. It asks her to predict outcomes in a cell-signalling pathway, explain her reasoning aloud, and compare her explanation with model answers, peer answers and a faculty rubric. It logs her misconceptions for the instructor, in a record she can inspect.
>
> The design follows the **Brynjolfsson**, **Li** and **Raymond** (2025) finding that AI gains accrue most strongly to less-experienced workers when AI scaffolds practice rather than replaces it; applied pedagogically, the same principle implies that AI is most useful as a tutor when it forces the learner to do cognitive work rather than receive cognitive product.
>
> In the laboratory, an AI assistant offers procedural prompts but is **forbidden** to draw conclusions; the conclusions must come from her data, her partner and her teaching assistant. The boundary is not a rule of conduct but a curricular design decision. It reflects the triadic principle that phenomenological observation and social validation cannot be delegated to functional agency.
>
> Mariana submits a process portfolio: prompts used, drafts written, errors caught, peer feedback, the lab report and a reflective note on what was delegated to AI and what was verified. Twice in the semester she defends her work in a short oral viva. The grade is composed: artefact, process and defence. The composition follows the

Bearman et al. (2024) principle that evaluative judgement, not the artefact alone, is what assessment under generative AI conditions must certify.

The pedagogical effect, accumulated over a term, is that Mariana learns to read AI outputs the way a literate reader learns to read newspapers — *sceptically, with provenance in mind*. The skill is not taught as a separate module on AI literacy. It is the by-product of an architecture in which AI is a participant under constraints that the student progressively internalises.

The architecture costs the institution more in teaching assistant time and instructor attention than the lecture-and-examination model it replaces. The reallocation is the substantive content of the curriculum redesign.

7.2 The History Seminar in Conversation with AI

In a third-year history seminar on twentieth-century decolonisation, AI is a participant, not a producer.

Students prompt the model for competing historiographical summaries of the period and bring them to the seminar. The pedagogical centre is not the output. It is the conversation: which archives are visible to the model, which voices are missing, what anachronisms or narrative conventions shape the answers, how dominant academic traditions are over-represented in training data, what counter-archives could correct the picture.

The model becomes a counterpart for the practice of source criticism — a function that historians have always performed against other sources and that now extends to a new kind of textual interlocutor.

The session is methodologically continuous with the long tradition of historical practice. The novelty is the explicit examination of the model as a representational artefact whose archive, training data and inference structure constitute their own historiographical positions. The seminar treats the model the way an earlier seminar might have treated a textbook, a colonial archive or a nationalist memorial: as *a constructed object whose claims must be interrogated*.

The pedagogical effect is to strengthen, not weaken, the student's capacity for source criticism. The framework of **Meira** et al. (2026) is deployed implicitly: the model contributes functional inferences; the seminar contributes social validation; the student contributes the integrative judgement that holds the historiographical position together.

The final assessment is an oral examination on a paired text and a model-generated commentary, *in which the student must distinguish what the model captured, what it distorted and what only a reader who has done the reading can say*. The historiographical imagination is **reinforced**, not replaced. The course publishes a description of its AI policy alongside its syllabus, in a form that any external examiner can **interrogate**. The publication is a small but consequential governance act: *it makes the pedagogical architecture inspectable and therefore defensible*.

7.3 The Business-School Capstone

A team of four MBA students is given a regional supply-chain problem proposed by a local employer.

AI agents support literature scans, market-data cleaning, scenario generation, financial modelling and presentation drafts. The team is required to maintain a **decision log**: what was delegated to AI, what was verified by which human, which assumptions were tested and which were not, where the model was wrong and how that was caught. *The decision log is the principal pedagogical instrument*. It externalises the meta-cognitive work that **Meira**, **Neves** and **Braga** (2025) identify as the defining competence of orchestrated cognitive work, and it makes that competence assessable.

The platform analysis in **Meira** (2025a) provides the conceptual frame for what the team is being asked to learn. Value in platform-orchestrated environments is generated by **composing** distributed capabilities **under** coherent governance. The capstone simulates the same task at the level of a single project: the team composes individual analysis, group deliberation, vendor expertise (the AI agents) and external stakeholder input (the employer) **under** a governance regime they had to design and have to defend. The deliverable **is no**t a recommendation. It is **a** recommendation, **a** process trace, and **a** defence of the orchestration choices that produced it.

The final deliverable is presented to a panel of faculty and the employer, with a live question session in which any team member may be asked any question. Grading **separates** the artefact from the process from the defence. The capstone treats AI as an enhancer of leverage and a generator of false confidence and assesses students on their capacity to use the leverage and resist the false confidence in equal measure.

The pedagogical model echoes the alignment principle of **Boud** and **Falchikov** (2006): *the capacity the institution wants to certify is also the capacity the*

assessment requires. The capacity the institution wants to certify **is** meta-cognitive orchestration; the assessment **requires** the team to orchestrate meta-cognitively under live, supervised conditions.

7.4 The Medical Clerkship

A clinical clerkship in a teaching hospital uses AI in three roles.

It pre-summarises patient records subject to faculty review; it offers differential-diagnosis prompts the student must accept, reject or extend with rationale; and it simulates rare-disease cases the student would not otherwise see.

Bedside teaching, however, remains the heart of the rotation. The clerkship explicitly **distinguishes** evaluable bedside skills — history-taking, examination, communication, ethical judgement, team practice — **from** AI-assistable cognition. The distinction follows the triadic logic of **Meira** et al. (2026): phenomenological observation of a patient and the relational competence of clinical communication **cannot be delegated** to functional agency, regardless of how capable that agency becomes.

Assessment includes mini-CEX observations, structured oral examinations, simulated emergencies and a clinical reasoning portfolio. Students who use AI to short-circuit reasoning are **not** punished by detection; they are **caught** by the structure of evaluation. The detection-based assessment regime that proliferated in 2023–2025 has been substantially abandoned in **serious** medical education because it **does not** work and because it teaches the **wrong** lesson; the structural alternative — *assessment so designed that AI cannot do the work for the student — is the only **durable** answer*. The shift is consistent with the **Bearman** et al. (2024) argument for evaluative judgement as the developmental target and with the **Boud** and **Falchikov** (2006) alignment principle.

The clerkship publishes its AI policy to patients and asks for consent at appropriate points, treating AI in care as **part of** the consenting relationship **rather than** as invisible infrastructure. The policy is informed by the high-risk classification of education and vocational training uses under the European AI Act (**European Parliament and Council**, 2024) and by the **NIST** (2023, 2024) AI Risk Management Framework's emphasis on documentation, transparency and human oversight. The clerkship therefore **models**, at the level of a single rotation, the governance **posture** the institution as a whole must adopt: *AI as a documented, supervised, contestable presence in the formative environment, not as an undeclared layer below the surface of professional practice*.

7.5 The Legal-Writing Studio

A second-year legal-writing studio uses AI as both a draftsperson and a foil.

> Students draft a brief without AI; produce an AI version with a documented prompt strategy; mark differences; identify model errors of law, citation, jurisdiction and tone; and rewrite a final draft. The instructor grades the rewrite and the meta-document, not the AI artefact. A short oral defence with adversarial questioning checks whether the student understands what she filed. The architecture is calibrated to the **specific failure modes** documented in recent legal practice: fabricated case citations produced by generative models that, in some American and European jurisdictions, have led to formal sanctions against the lawyers who filed them.
>
> The studio runs an integrity unit in which such fabricated citations are dissected, traced and corrected. Students leave with a working **method** for using AI in practice and a credible defence of their own competence to bar examiners, supervising partners and clients.
>
> The professional stake is direct. Under the **Autor** and **Thompson** (2025) framing, AI **compresses** the routine expertise that junior associates historically **learned** by performing and raises the premium on integrative judgement: the capacity to know what the brief must accomplish, what authorities support it, what risks attend its filing and how to defend it under interrogation. The studio is **designed to develop** that integrative judgement in the same setting in which AI use is normalised.
>
> The studio also illustrates the pedagogical use of **process-visible work**. The student's sequence — clean draft, AI version, comparison, rewrite, defence — produces a **trail** that the instructor can grade for **orchestration competence** as well as for the artefact itself. The trail **externalises** the meta-cognitive work that **Meira**, **Neves** and **Braga** (2025) identify as the new professional competence. The instructor's craft is the **design** of the sequence; the student's craft is the disciplined **operation** inside it.

7.6 The Vocational Adult Learner on a Mobile-First Pathway

Antônio is a forty-two-year-old electrician retraining as an industrial automation technician through a state public university extension programme in northeastern Brazil.

> He has a smartphone and intermittent broadband. The programme is mobile-first and multilingual. An AI tutor in his first language explains concepts, generates practice problems calibrated to his progress, and translates technical

documentation from manufacturers whose manuals are predominantly in English. A local mentor — a senior technician at a partner firm — meets him weekly for supervised practice on real equipment. Recognition of prior learning credits his decades of trade experience under criteria developed with the regional employers' association.

The programme illustrates how the **Council of the European Union**'s (2022) Recommendation on micro-credentials translates into a Global South context. Antônio earns **stackable** credentials linked to a national qualifications framework, each defined by demonstrated competence rather than seat-time. AI **does not** replace his mentor; it makes the mentor's time count, by handling explanation and translation between the mentor sessions and the practice exercises.

The architecture is deliberately designed **not to** imitate an elite residential campus. It is designed to **honour** the apprenticeship logic and lifelong learning needs of a society that needs many more skilled technicians **than it has room for** in conventional university classrooms.

The triadic composition is the conceptual core. Antônio's individual intelligence — embodied competence, accumulated tacit knowledge, motivated learning — is the **ground**. The social intelligence of the mentor, the cohort of fellow learners and the partner firm is the **validating environment**. The artificial intelligence of the tutor and the translation system is the **scaling layer**. The architecture is the operational form of the triadic synergy that **Meira** et al. (2026) identifies in formal terms. It is also a concrete answer to the apprenticeship paradox of section 4.2: the apprenticeship is **reconstituted** inside the educational architecture, with the mentor relationship **deliberately resourced** rather than displaced by automation.

7.7 The Global-South Community-Learning Hub

A small public university in a low-income region operates a community-learning hub in a town where most residents have neither stable internet nor a home computer.

The hub provides shared devices, a librarian, an academic counsellor and a low-bandwidth AI tutor that can run with intermittent connectivity. The curriculum is locally **co-designed** with health-service, agricultural-extension and municipal partners. Credentials are stackable, articulated with regional employers, and recognised across the state system. The AI is configured to operate in regional Portuguese and at least one Indigenous language; data are stored under national jurisdiction; the model is auditable; and the institution exercises refusal as a

governance capacity for uses (such as automated discipline or unaccountable scoring) it deems incompatible with public mission.

The hub illustrates what AI-native looks like at the **periphery** of the global system — not a deficit **imitation** of the centre but an institutional **design** that the centre will eventually need to learn from. The phygital architecture that underwrites such hubs is the subject of **Meira** and **Meira** (2026). Its triadic ontology — Physical (the building, the devices, the bodies present), Networked Digital (the AI, the data, the credentials), Networked Social (the partner organisations, the staff, the validated community of practice) — gives **explicit form** to the institutional architecture that an isolated platform deployment would obscure. The Physical dimension is foundational: the hub exists **because** the building exists, **because** the devices are present, **because** the librarian and counsellor are physically there to greet learners.

The hub also illustrates the political-economy stake of section 11. Without public investment in physical and digital infrastructure, multilingual AI development, professional pathways for the staff and recognition mechanisms across institutional boundaries, the hub **cannot exist**. *The institutional architecture is therefore a policy artefact as well as a pedagogical design*. **UNESCO**'s 2026 framing of higher education as a public mission applies in concrete form: the hub is the operational expression of the mission at scale where the elite campus cannot reach.

7.8 The Faculty Workday: From Grading to Orchestration

Professor Silvia coordinates a high-enrolment first-year course.

Her workday two years into AI-native reform looks different than it did. She no longer drafts model solutions; the AI does. She no longer answers the same procedural email forty times; the AI does, under her supervision. She **does**, however, spend more time on three things: **designing** assessment evidence that AI cannot generate, **supervising** teaching assistants who orchestrate AI-mediated learning, and **meeting** individual students whose **process logs** indicate confusion. Her promotion file now includes evidence of teaching reform, assessment validity, supervisory work, scholarly contribution and academic-work redesign — **not only** research publications.

The workload model under which she now works was negotiated with the union, the dean and the provost. It explicitly recognises **design**, **supervision** and **assessment** alongside teaching delivery and research. *Her workload did not decrease in hours; it relocated in content*. The relocation was visible, contested and bargained, which is what made it sustainable.

The contrast with the silent reorganisation analysed in **Meira**, **Neves** and **Braga** (2025) under the heading of digital Taylorism is deliberate. *The digital-Taylorist trajectory uses AI to fragment academic judgement and reroute it through dashboards; the orchestration trajectory uses AI to relocate academic judgement into higher-leverage activity*. The institutional difference between the two trajectories is **governance**: which actors participate in the redesign, on what terms, with what protections.

Professor Silvia's working day also illustrates the reflexive condition of section 4.3. The institution is asking students to **become** meta-cognitive orchestrators of their own learning under AI conditions.

Professor Silvia is the institutional embodiment of the same competence. She **models**, in her own working practice, the orchestration the institution claims to certify. The reflexive condition is not a moral expectation. It is a structural requirement: a faculty whose own work is dashboard-mediated drudgery cannot credibly form students into reflective orchestrators of cognition. The orchestration trajectory is therefore the only one consistent with the institution's educational claim.

7.9 Design Patterns Derived From the Vignettes

The vignettes share a structure.

AI is most educationally valuable when it is **embedded** in a sequence — preparation, practice, feedback, collaboration, human judgement, performance and reflection — rather than substituted for a step in that sequence.

Assessment becomes credible when it **combines** artefact, process and defence, in line with the **Bearman** et al. (2024) analysis of evaluative judgement under generative AI conditions.

Faculty remain **essential** as designers of the sequence, judges of performance and stewards of academic standards.

Students gain **capability** when they are required to direct, evaluate and defend AI use, rather than to consume or hide it.

Equity is **not** a downstream effect: it is the choice of where to invest human time.

Three further patterns recur across the vignettes.

First, the boundary between AI-assistable and AI-resistant work is decided **pedagogically**, not technologically. The clerkship's decision that bedside competence is not AI-assistable is not a claim about model capability; it is a curricular and ethical decision about what professional formation requires.

Second, the **process trail** — decision log, process portfolio, reflective note, public AI policy — **is the principal pedagogical innovation**. It externalises the meta-cognitive work and makes it assessable.

Third, the *institutional architecture is as much a design object as the pedagogical sequence*. Workload models, professional development, mentor relationships, partner agreements and governance documents are part of the curriculum, not its background. Table 2 distils the recurring design patterns.

Table 2. AI-native design patterns for learning and assessment

Design pattern	Core idea	Suitable uses
AI-supported preparation	Students use AI for explanations, retrieval practice and misconception detection before human contact.	Large introductory courses; foundation modules; bridging programmes.
Process-visible work	Students submit prompts, drafts, decision logs, source trails or model critiques alongside the final artefact.	Writing, design, coding, research methods, professional projects.
Deliberate AI-free intervals	Students perform selected tasks without AI to demonstrate foundational fluency.	Core reasoning, mathematics, clinical knowledge, language proficiency.
Authentic task with AI governance	AI is allowed, but students must justify delegation, verification and judgement.	Capstones, studios, clinics, consulting projects, engineering design.
Oral or live defence	Students explain, defend and extend their work in real time.	High-stakes certification; honours projects; professional programmes.
Portfolio with longitudinal evidence	Competence is demonstrated through accumulated artefacts and reflections.	Modular credentials, lifelong learning, interdisciplinary programmes.
Simulation and scenario testing	Students act under changing conditions and receive feedback.	Medicine, law, education, business, public administration, crisis response.
Community or workplace problem lab	Learning is organised around external problems with public or employer partners.	Regional development, applied research, civic education, reskilling.

Reflexive academic-work redesign	Academic teams map teaching, assessment, advising and administrative work across individual, social and artificial intelligences.	Faculty workload redesign; teaching-assistant preparation; student services; assessment reform.

8. Compositional Transformation: Institutions, Organisations, Ecosystems

8.1 Three Layers That Must Move Together

If higher education is a hyperobject, the unit of reform **cannot be** the isolated classroom, course, platform or degree. It must be the composition of institutional rules, organisational architectures and ecosystem arrangements. The distinction adapts **North**'s (1990) classical separation between institutions as the rules of the game and organisations as the players, with the ecosystem as the surrounding field of relations. The conceptual move matters because it locates reform in the **right** place.

> The classroom is not the institution; *the institution is the set of rules within which the classroom is possible, organised and evaluated*.

The **institutional layer** includes law, accreditation, funding, quality assurance, data protection, academic-labour regulation, professional licensure, credential recognition and AI governance. It is slow but it defines what organisational forms are legitimate. The **Council of the European Union**'s (2022) Recommendation on micro-credentials, the European AI Act (**European Parliament and Council**, 2024), the **NIST** (2023, 2024) AI Risk Management Framework, national accreditation regimes and the various national AI policy frameworks under development belong to this layer. The institutional layer changes through legislative and regulatory processes whose timescales are typically measured in years and whose direction is the consolidated outcome of contested political work.

The **organisational layer** includes the internal design of universities and providers — curriculum, assessment, faculty roles, student support, data infrastructure, pricing, partnerships, research strategy and governance. It is where inherited routines are strongest, because organisations are the historical embodiment of the rules of the game and embody them in workload models, promotion criteria, procurement processes, governance committees and professional norms. **Clark**'s (1983) comparative analysis of academic organisations across national systems maps the structural diversity of this layer. **Slaughter** and **Rhoades**'s (2004) academic capitalism describes its present political-

economy condition. AI-native organisational reform must work within these structural realities, not around them.

The **ecosystem layer** includes education markets and the sectors served by education — employers, professional bodies, governments, families, learners and other educators. The ecosystem is where credentials acquire (or lose) value, where research circulates, where partnerships are formed and where political legitimacy is granted or withheld. The **Boston Consulting Group**'s analysis (2026) and the Deloitte trends report (**Clark** et al., 2026) describe ecosystem signals at the level of demand, employer practice and partnership behaviour; **UPCEA** (2025) reports the same signals from inside the modular and adult-learning ecosystem.

8.2 Why One-Layer Reforms Fail

Single-layer reforms produce **reliable** disappointments. A university cannot redesign assessment if accreditation still demands inputs and seat-time. A regulator cannot recognise modular credentials if employers do not trust them. An employer coalition cannot rebuild the entry contract if universities cannot certify what the coalition values. A faculty group cannot redesign assessment if workload models do not recognise the **new work**. A procurement office cannot impose audit rights on vendors if the institution does not have the legal infrastructure to enforce them. Each example is a familiar pattern in higher education reform; each example is a layer **misalignment**.

The pattern has a structural explanation. Each layer encodes the rationality of its actors at the timescale on which **they** operate. Regulators encode the rationality of their political principals at a multi-year timescale. Organisations encode the rationality of their internal coalitions at a multi-year timescale. Ecosystems encode the rationality of distributed actors at a continuous timescale. When the layers are **misaligned**, transformation oscillates: organisations innovate, regulators sanction or fail to recognise the innovation, ecosystems sceptically wait, innovation stalls, and the next round of reform begins from the same misalignment.

> The platform analysis in **Meira** (2025a) offers the same conclusion in different language: value migrates to whichever layer is **best aligned** at the moment of disruption, and incumbents that fail to align lose the ability to set terms.

The **compositional principle**, then, *is not that reform must move on all fronts at once, but that no front can move sustainably without anticipating where the other fronts must move next*.

> The institutional task is **to design reform sequences** in which each step is supported by the next, and **to maintain enough conversational and political**

infrastructure across layers that misalignments are identified before they harden. The implementation architecture proposed in section 9 is structured to make this multi-layer sequencing explicit at four time horizons. The compositional framework presented in Table 3 below specifies the alignment requirements at each layer.

Table 3. Compositional transformation framework for AI-native higher education

Layer	Focus	Transformations
Institutions	Rules, norms and regulatory models	Recognition of modular credentials; AI-specific quality assurance; learner data rights; accreditation reform; professional licensure update; AI Act-aligned governance.
Organisations	Methods, processes and business models	Versioned curricula; human-AI pedagogical design; authentic assessment; AI-supported advising; new faculty workload models; redesigned business models for human attention.
Ecosystems	Education markets and markets served by education	Credential marketplaces with quality signals; employer-integrated pathways; public AI infrastructure; civic and research partnerships; shared assessment infrastructure.

8.3 Layer Interactions and Feedback Loops

Layer interactions produce feedback loops that determine the speed and equity of reform. Regulatory recognition **enables** organisational experimentation. Organisational credibility **shapes** ecosystem trust. Ecosystem demand **pressures** institutions to update rules. The loops can run in virtuous or vicious directions. A virtuous loop runs when public investment in AI infrastructure **enables** organisational redesign, which **produces** credentialed outcomes that ecosystems trust, which **generates** political support for further regulatory progress. A vicious loop runs when platform capture **undermines** organisational autonomy, which **degrades** the quality of credentials, which **erodes** ecosystem trust, which **generates** political pressure for restrictive regulation that further constrains organisational autonomy.

Which loop runs depends on the **choices** made at each layer in the present period. The **Boston Consulting Group**'s analysis (2026) describes the upside loop; the platform-capture analysis in **Meira** (2025a) and the digital-Taylorism analysis in **Meira**, **Neves** and **Braga** (2025) describe the downside loop. The two loops are **not** equally likely under all institutional conditions. Public investment, multi-stakeholder governance, learner rights, faculty co-design and refusal as a governance capacity (section 13) tilt the system toward

the virtuous loop. Market deference, vendor dependence, unilateral organisational restructuring and the marginalisation of labour voice tilt the system toward the vicious loop.

The compositional principle also has a **temporal** dimension. *Sequencing matters as much as content*.

> If organisational reform runs ahead of institutional recognition, *the reform exhausts itself in compliance work*.
>
> If institutional recognition runs ahead of organisational capacity, *the recognition produces credentials no one trusts*.
>
> If ecosystem demand runs ahead of organisational supply, *the demand is met by platform providers operating outside academic governance*.

The implementation architecture proposed in section 9 is sequenced to keep the three layers in conversation. The political economy of that sequencing — who participates in the conversation, on what terms — is the substantive content of section 11.

9. Implementation: What Leaders Do Across Four Horizons

9.1 First Principles for Sequencing

The most common institutional failure is not lack of awareness. **It is lack of sequencing**.

> Leaders can accept the **need** for AI-native transformation **while** being unable to answer the practical question: *where do we start without breaking the institution*?

The answer is not a universal recipe. It is a sequencing logic. **Three** principles structure that logic.

> **First, begin where trust is most fragile**; in most institutions, that is assessment. If assessment cannot credibly certify individual and collective competence under AI conditions, every other claim about learning weakens, as the Bearman et al. and **Boud** and **Falchikov** (2006) literature makes clear.
>
> **Second, do not confuse policy statements with institutional capacity**; staff time, support units, professional development and governance are the fundamental binding constraints.
>
> **Third, redesign academic work in parallel with student-facing reforms**; otherwise the reflexive contradiction analysed in section 4.3 will silently undermine every other change.

A **fourth** principle follows from the red-swan analysis of **Meira** (2025b):

> *sequence under conditions of irreducible uncertainty by building option value rather than by optimising a single trajectory*.
>
> The implementation plan must be robust to scenarios the institution cannot now foresee — the labour-market reconfiguration that **Massenkoff** and **McCrory** (2026) document at the entry level, the platform-capture pressures analysed by **Meira** (2025a), the regulatory shifts hinted at by the European AI Act and its emerging national counterparts.

The implementation architecture below is structured to keep multiple paths open rather than to commit to a single transformation pathway whose premises may be invalidated by external rupture.

The fifth principle is participation. The political economy analysed in section 11 implies that **any** sequencing decision made **unilaterally** by senior leadership will encounter resistance **proportional** to the unilateralism.

> Sequencing decisions must be made with the actors who will live them: faculty, students, professional staff, unions where they exist, employer partners, regulators and the public mission constituencies the institution serves. Participation is not a tax on speed; it is the condition of durable reform. With those principles in hand, the next four subsections offer a navigable pathway across four horizons.

9.2 The First Ninety Days

In the first ninety days, leaders should establish an **AI education governance group** with real authority and broad representation:

> academic leadership, faculty, students, professional staff, information security, library, accessibility, legal counsel, data protection, teaching and learning experts, ethics representation and external advisors.

The group should be a standing body, **not** a committee assembled for the AI moment and then dissolved; the conditions it addresses **will not** be settled in any single planning cycle. Its authority **must** extend to recommending policy, reviewing procurement, approving prototypes and commissioning the institutional inventory.

The group commissions **four** things. **First**, an inventory of existing AI uses across the institution — formal procurements, departmental experiments, individual faculty practice and student-led use. **Second**, a rapid audit of assessment exposure, programme by programme, identifying which assessments can be substantially produced by AI without detection and which therefore require redesign or supplementation. **Third**, a procurement

review of existing contracts for audit rights, interoperability, data-protection guarantees, accessibility and exit clauses, against the framework articulated in the **NIST** (2023, 2024) AI Risk Management Framework and its Generative AI Profile (**NIST**, 2023, 2024) and the EU AI Act (**European Parliament and Council**, 2024). **Fourth**, a published interim guidance for students and staff that is specific enough to be useful, provisional enough to be revisable, and produced with **their** participation.

The group also chooses a **small portfolio** of prototypes — three to five — that cross the whole learning architecture rather than concentrating on a single tool deployment. The vignettes in section 7 illustrate the kind of prototypes that test multiple dimensions of the framework simultaneously: a redesigned high-enrolment first-year course, a capstone with explicit AI orchestration, a clinical or studio rotation with process-visible assessment, an adult-learning pathway with stackable credentials, a faculty-workday redesign in one department. Each prototype is a learning instrument, not a pilot to be scaled blindly.

The Monday-morning agenda is therefore concrete: **convene** governance with authority; **audit** existing AI uses; **stress-test** assessment; **choose** prototypes that cross the whole architecture; **invest** in faculty and professional staff; **publish** student guidance that is specific, fair and reviewable.

9.3 Twelve Months

Within twelve months, the institution should have a **distributed AI literacy requirement** embedded in programmes across the disciplines, not concentrated in a single elective; **assessment-redesign exemplars** in multiple disciplines that other programmes can adapt; **staff development programmes** with workload recognition, designed for the realistic time budgets of faculty and professional staff rather than as additional uncompensated work; and a **procurement framework** requiring auditability, interoperability, accessibility and data protection.

> The framework should explicitly reference the EU AI Act's risk classification for educational uses and the **NIST** (2023, 2024) framework's vocabulary, so that the institution's procurement language is interoperable with the emerging international regulatory landscape.

The institution should also have a **student-facing AI guidance system** that is updated transparently and produced with student participation; a **data-rights statement** for learners that specifies what records the institution holds, what rights students have over them and how those rights can be exercised; a **research stream on the effects of AI** on learning, internal or external, that feeds into institutional decision-making; and a **teaching-and-learning support unit** large enough to make assessment redesign feasible at scale.

The size of the support unit is the binding constraint at this horizon. Without sufficient professional capacity in instructional design, assessment expertise and pedagogical research, the institution's policy statements will **outrun** its operational ability.

Implementation at this horizon should be visibly under-budgeted in money and over-invested in staff time, because the *binding constraint is human capacity*, not licensing. Procurements that promise to substitute for staff capacity should be received with **scepticism**; the digital-Taylorism risk analysed in **Meira**, **Neves** and **Braga** (2025) is **most acute** at exactly this horizon, when fiscal pressure and AI enthusiasm combine to make capital substitution for labour an **attractive** narrative. The countermove is to *commit explicit human-time budgets to assessment redesign, faculty development, professional-staff development and governance*, and to defend those budgets against the recurrent pressure to convert them into licensing spend.

9.4 Three Years

Within three years, the institution should **move from pilots to architecture**. Programmes should be mapped against graduate capabilities for AI-mediated work and citizenship; the mapping should be public and revisable. High-stakes assessment should include at least one **credible** form of supervised, oral, practical, simulated or process-visible evidence in **every** programme, following the design patterns of section 7 and the framework articulated by Bearman and colleagues (2024). The library and learning-design functions should be **reorganised** as an academic-work redesign service: the previous division between technological support, library and teaching expertise breaks down under AI conditions and the reorganisation should **reflect** that breakdown.

Faculty workload models should explicitly **value** design, supervision, mentoring and assessment, with promotion criteria updated accordingly. The reform **must** include adjunct, graduate-student and professional-staff labour, not only tenured faculty; **otherwise** it will reproduce the stratification it claims to address. The political-economy analysis of section 11 is operationally relevant here: workload reform that does not include all relevant labour constituencies **will be** resisted, eroded or quietly reversed.

> Procurement should **require** interoperability, audit rights, accessibility, security, data minimisation and **exit**; *the institution should be able to leave any vendor on documented terms without losing its students' records or the integrity of its credentials*.

Modular credentials should be embedded in degree pathways, not bolted on. The **Council of the European Union**'s (2022) Recommendation provides one architecture; national

qualifications frameworks provide others. *The principle is that the credential travels with the learner and stacks coherently into formation*.

> *The institution should publish, at this horizon, an honest annual statement of what its graduates can and cannot do with and without AI, how those claims are assessed, how student data are governed and how faculty work has changed*.
>
> **The statement is a governance artefact**: it makes the institution's claims inspectable, contestable and improvable. The **Boston Consulting Group**'s analysis (2026) treats this kind of public statement as a signal of credibility in the AI-native market; under the political-economy reading of section 11 it is also a discipline against drift.

9.5 Ten Years

By 2035, AI-native transformation should be **visible** in the institution's public claims.

> *The institution should be able to say, with evidence, what its graduates can do, how those claims are assessed, how student data are governed, how faculty work has changed, how human support is protected, how credentials are recognised, how research is conducted under AI conditions and how the institution contributes to public missions*.

The decade horizon is not a forecast; it is a **commitment** device. It specifies what the institution holds itself accountable to demonstrate at that point. Institutions that **cannot** make such claims will, by 2035, occupy the **residual** market between platform providers and elite incumbents — a position that the platform analysis in **Meira** (2025a) identifies as structurally **unsustainable**.

At the ten-year horizon, the **maturity model** in Table 4 below becomes a strategic instrument. Institutions can locate themselves on the model, identify the level their mission and context require, and chart the alignment work between the present level and the target level. The model is **not** a ranking tool. Its purpose is to support strategic **conversation** across institutions that are not directly comparable in mission, scale, resources or context. A regional public university and a research-intensive global university may both legitimately be at Level 3 (Redesign), pursuing different forms of the same alignment requirement.

The ten-year horizon also brings into view the deeper **reflexive** condition. By 2035, the institution's **own** academic and professional labour should be visibly **organised** around the same orchestration logic the institution asks its graduates to internalise. Workload models, promotion criteria, professional development and governance structures should reflect what **Meira**, **Neves** and **Braga** (2025) analyse as **meta-cognitive orchestration** at the level

of the **institution's own work**. The reflexive test is what makes the institution's AI-native claim credible at a generation's distance.

9.6 A Maturity Model for Self-Study

The four horizons can be paired with a seven-level maturity model that institutions can use for self-study. The levels are descriptive rather than prescriptive: an institution may legitimately occupy different levels in different programmes, and the trajectory across the levels is not linear. The model is presented in Table 4. It is not a ranking tool. Its purpose is to support strategic conversation, including across institutions that are not directly comparable, and to make explicit the kinds of risk associated with each level. The level descriptions intentionally include the dominant failure mode, so that institutions can plan against the failures characteristic of their current condition rather than only against the failures of the level they aspire to.

Table 4. AI-native higher education maturity model

Level	Characteristic posture	Evidence of practice	Main risk
0. Denial	AI is treated mainly as misconduct or distraction.	Bans, detection tools, inconsistent guidance.	Loss of credibility as student practice and workplace practice diverge.
1. Enablement	AI tools are adopted in existing processes.	Chatbots, pilots, faculty experimentation, generic policies.	Fragmentation, inequity and vendor dependency.
2. Alignment	AI use is linked to institutional priorities.	Governance group, procurement rules, staff development, programme pilots.	Strategy remains rhetorical if assessment and workload do not change.
3. Redesign	Core academic architectures are rebuilt.	AI literacy across curriculum, authentic assessment, data rights, redesigned faculty roles.	Implementation fatigue and uneven quality across departments.

4. Ecosystem integration	Institution participates in shared public and professional infrastructures.	Portable credentials, employer pathways, independent assessment, public AI services, lifelong learning.	Capture by platform markets or narrow employability metrics.
5. Public renewal	AI-native design strengthens public mission and trust.	Demonstrable learning, equity gains, civic contribution, research capacity, accountable governance.	Complacency or prestige insulation from continued change.
6. Reflexive renewal	The university redesigns its own work as an explicit site of meta-cognitive orchestration.	Workload models, promotion criteria, staff roles, governance processes and support units reflect the same logic asked of students.	Internal stratification if orchestration roles accrue to elites while routine monitoring and platform labour are pushed downward.

9.7 The Monday-Morning Agenda

A provost or rector can sequence the next ninety days against this model. The aim is not to ascend the levels **mechanically**; it is to identify where the institution actually is, what level its mission, demographics and capacity require, and which dimensions of Table 1 must be aligned to move credibly. The maturity is **not** a virtue but a question: maturity for what, for whom and at what cost? An institution serving a working-class urban population may need to be at Level 4 (Ecosystem integration) to meet its mission; an elite research-intensive institution may legitimately operate at Level 5 (Public renewal) by virtue of its research and public engagement role; a small liberal-arts college may organise its strategy around Level 6 (Reflexive renewal) at the level of its faculty culture.

The Monday-morning agenda has a concrete form regardless of institutional type.

> Convene governance with authority. Audit existing AI uses. Stress-test assessment. Choose prototypes that cross the whole architecture. Invest in faculty and professional staff. Publish student guidance that is specific, fair and reviewable. Open the procurement file. Talk to the union, if one exists, before policies are drafted. Talk to the accreditor about what kinds of evidence will count. Talk to local

employers about what they would actually trust. Talk to students about what they are actually doing.

Each of these conversations is a small **institutional** act. Cumulatively, over ninety days, they reposition the institution **from** a defensive **to** a designing posture. The compositional framework of section 8 ensures that the design touches the layers that matter rather than only the ones most visible to senior leadership.

10. Policy Architecture: Principles, Instruments and Safeguards

10.1 Purpose and Ten Policy Principles

AI-native higher education requires more than institutional entrepreneurship. It requires policy architecture. Governments, quality-assurance agencies, accreditation bodies, professional regulators and funders shape what institutions can responsibly attempt.

Policy should neither freeze inherited forms nor accelerate uncritical adoption. Its purpose can be formulated as follows:

> *to develop, throughout life and across society, the **human** capacity to **understand**, **create** and **act with, about** and **despite** artificial intelligence, while advancing knowledge, public reason, social mobility, democratic agency and planetary responsibility*.

The formulation is normative and intentionally so. It is the **political** and **ethical** orientation against which specific instruments are designed and evaluated.

The formulation builds on **Marginson**'s (2016) recovery of higher education as **common good** and on **Barnett**'s (2011) philosophical defence of *being a university* as an institutional commitment to critical inquiry. It is also consistent with **UNESCO**'s 2026 framing of higher education as a public mission facing technological, geopolitical, demographic and funding pressures (**UNESCO**, 2026). The purpose is **not** jurisdiction-specific. Different national systems will translate it through different instruments — public funding architectures, accreditation regimes, professional licensure, data-protection law, AI governance — but the underlying commitment travels across systems.

This purpose implies ten policy principles.

> **First**, learners are subjects of education, not data objects; they should have meaningful rights over learning records, credential evidence and AI interaction data, and those rights should be enforceable.
>
> **Second**, demonstrated competence matters more than time served, while protecting coherent formation against fragmentation.

Third, modular credentials require coherent quality assurance and cannot be allowed to proliferate as an opaque market.

Fourth, academic and professional labour must be redesigned rather than merely re-tooled; the digital-Taylorism trajectory analysed in **Meira**, **Neves** and **Braga** (2025) is the principal labour-market risk policy must constrain.

Fifth, public infrastructures must coexist with private platforms; the platform analysis in **Meira** (2025a) indicates that without deliberate public capacity, orchestration value migrates exclusively to private actors.

Sixth, equity is a design constraint, not an afterthought; AI-native systems must be evaluated by their effects on access, persistence, learning quality, outcomes and support for disadvantaged learners.

Seventh, lifelong learning becomes a core function rather than a peripheral revenue stream; the recognition of prior and informal learning is a structural requirement, not a marketing feature.

Eighth, accreditation evidence shifts from inputs to learning, equity, governance and public value, with active safeguards against compliance gaming.

Ninth, AI governance in education must be embedded in human-rights, data-protection and educational frameworks, including the high-risk classification of education and vocational training uses in the EU AI Act (**European Parliament and Council**, 2024) and the **NIST** (2023, 2024) risk-management vocabulary (**NIST**, 2023, 2024).

Tenth, transformation must be reflexive: applied to the institution's own labour as the operational test of its educational claim.

10.2 Instruments: Credentials, Data, Procurement, Licensure

Policy instruments follow from these principles. Governments and quality-assurance agencies should recognise **modular**, **stackable** and **portable** credentials under rigorous standards for assessment, coherence and learner protection. The European Council's Recommendation on micro-credentials (**Council** of the **European Union**, 2022) offers one institutional template; comparable frameworks are emerging in other systems. The instrument is **not** modularity for its own sake; it is the **recognition** that coherent formation can be assembled across providers under shared quality standards. Without such recognition, modularity becomes badge **inflation**; with it, modularity becomes a **credible** alternative to the monolithic degree.

Data legislation and sector regulation should define learner rights over educational records and AI-generated traces, including rights to access, port, correct, contest and where appropriate delete. AI systems used in education may directly affect access, progression and evaluation; the EU AI Act treats several education and vocational training uses as high-risk and therefore subject to documentation, oversight, transparency and conformity obligations. The implementing detail of the Act will be litigated and revised over the coming years; the policy task for higher education is to keep its instruments compatible with the architectural intent of the Act rather than to wait for definitive guidance.

> Public procurement rules **should require** audit rights, interoperability, accessibility, security, data minimisation and exit options. Public institutions **should not** outsource core educational standards to opaque models or vendor-controlled analytics; the procurement instrument is what makes this commitment operational.

The **NIST** (2023, 2024) AI Risk Management Framework and its Generative AI Profile (**NIST**, 2023, 2024) provide a vocabulary that procurement offices can deploy without binding themselves to specific vendors. The instrument has two faces. It is a defensive measure against vendor lock-in and platform capture. It is also a **developmental** measure: by aggregating procurement standards across institutions, public systems shape the market the vendors must serve.

Professional bodies should **revise** licensure pathways to recognise AI-mediated work while protecting public-interest standards in medicine, law, engineering, accounting and other regulated professions. The pathway revision is **not** a deregulation. It is a recalibration of what the licensure must certify: integrative judgement, accountability for AI-assisted decisions, capacity to supervise AI under professional ethics. The **Bearman** et al. (2024) analysis of evaluative judgement is directly relevant; the licensure pathway is the institutional form in which evaluative judgement is certified at the professional level. Without such recalibration, the **gap** between what professionals do and what licensure certifies will **widen**, and the licensure will progressively **lose** public trust.

10.3 The Dual Risk of Outcomes-Based Assurance

The shift from inputs to outcomes is correct in theory but historically perilous in practice.

> The history of outcomes-based quality assurance is littered with reductive metrics, compliance games, narrow audit cycles and regulatory capture — the very risks identified elsewhere in this paper.
>
> If AI-native accreditation **merely** replaces seat-time counting with dashboard counting, it will reproduce the failure mode under a new name.

> The institutional record of outcomes-based reform — across health, education and public administration in multiple jurisdictions over the past three decades — gives **no warrant** for optimism that the same instrument applied to AI-native higher education will avoid the same fate without explicit countermeasures.

The countermeasures are not exotic. They include **qualitative** evidence alongside **quantitative** indicators; **external** moderation and sampling to limit the gameability of self-reported data; **explicit** faculty professional **judgement** as a recognised source of evidence; learner **voice** in the assessment of the assessment regime; and the capacity to **refuse** reductive metrics when they distort the very outcomes they are intended to measure.

> The strongest outcomes regimes will be those that explicitly **limit** what counts as evidence and reserve a defended space for the unquantifiable. The principle is consistent with the institutional theory of **North** (1990): *institutions that confuse measurement with the underlying object they intend to govern produce dysfunctional reform*.

The policy architecture must therefore acknowledge that any outcomes regime requires the active **maintenance** of professional judgement against the recurrent pressure to convert **all** evidence into countable form. The maintenance is not a single design decision. It is the ongoing political work of regulators, accreditors, professional bodies and institutions to preserve the qualitative and judgement-based components of evaluation against the pressures of audit efficiency, comparability and dashboard visibility. The historical record of academic capitalism described by **Slaughter** and **Rhoades** (2004) suggests that without explicit countermeasures, the pressures **will** prevail.

10.4 Avoiding Compliance Theatre

Compliance theatre is the predictable failure mode of well-meant regulation. Universities will write policies they do not implement. Vendors will produce dashboards that do not measure what matters. Accreditors will inspect what is easily inspected. Faculty will sign off on assessment regimes whose effects they have not tested. Students will be presented with rights they cannot operationally exercise. The **pattern** is familiar across regulatory domains and there is no reason to expect AI-native higher education to be exempt unless explicit measures are taken.

The countermeasure **is not** more regulation **but** better-designed regulation:

> regulation that requires evidence of practice rather than evidence of intent;
>
> that includes independent learner and faculty voice;
>
> that funds the capacity to comply meaningfully;

that includes sunset clauses on rules whose value has expired;

and that exposes regulatory action to public scrutiny in a form that any interested actor can audit.

The AI Act's emphasis on documentation, transparency, post-market monitoring and conformity assessment, properly localised, points in this direction. The **NIST** (2023, 2024) framework's emphasis on practice-oriented risk management, similarly, asks for evidence rather than ritual. The instruments **exist**; the political work is to use them **seriously**.

The deeper countermeasure is institutional **culture**. A regulator that operates in genuine conversation with the regulated produces less compliance theatre than a regulator operating at a distance. A university whose senior leadership engages with the accreditor's actual concerns produces less compliance theatre than one optimising for the visible metric. A faculty culture that treats assessment redesign as a matter of academic integrity rather than as a compliance task produces less compliance theatre than one that does the opposite. *The cultural countermeasure cannot be legislated; it can be incentivised, modelled and protected, and policy design should attend to those incentives explicitly*.

11. The Political Economy of Change

11.1 Transformation as Political Process

AI-native transformation is not an optimisation problem. It is a political process.

> It redistributes resources, authority, labour, legitimacy and risk. It changes who designs learning, who evaluates competence, who owns data, which vendors gain power, how faculty time is valued, how employers influence credentials and how regulators measure quality.

The political economy must be addressed **directly**, not bracketed. Treating reform as a **managerial** exercise — a project plan to be executed by senior leadership through delegated teams — is the **most common mistake** in higher education change management and the **most reliable predictor of failure**.

The Schumpeterian language of **creative destruction** is useful here, but only if its political dimension is held in view. **Schumpeter**'s (1942) analysis identifies the recurrent process by which new economic combinations destroy the value of existing arrangements; *the political question that his analysis underweights is who absorbs the destruction and who captures the creation*.

In higher education, the **destruction** is concentrated in routine cognitive labour, in unsupervised written assignments, in inherited workload models, in vendor-locked infrastructures and in credentials whose signal has degraded.

The **creation** is concentrated in orchestration roles, authentic assessment infrastructures, modular credentials, public AI capacity and reflexive academic-work redesign.

Which constituencies bear the destruction and which capture the creation is the political question of the transition.

The platform analysis in **Meira** (2025a) adds a structural dimension.

> *Under platform-economy conditions, the distribution of created value tilts systematically toward whichever actor controls the orchestration layer of the value chain*.
>
> *If universities allow private platforms to orchestrate the educational layer of their work, the value created by AI-native transformation will accrue disproportionately to those platforms regardless of where the transformation is operationally executed*.

Defending institutional control of the educational orchestration layer is therefore **not** a defensive posture; *it is a political-economy precondition for the institutional capture of the value the transformation creates*.

11.2 Internal Arenas

Inside universities, change encounters multiple veto points.

Departments may resist common AI literacy requirements if they perceive threats to disciplinary autonomy.

Faculty may resist assessment redesign if it increases workload without recognition.

Academic senates may defend inherited forms of evaluation.

Adjunct and graduate-student labour may be silently restructured under the banner of efficiency, with the predictable downward pressure on precarious labour that **Slaughter** and **Rhoades** (2004) document in their analysis of academic capitalism.

Disciplinary prestige hierarchies may distort which programmes attract investment and which absorb risk.

None of this is illegitimate; *much of it is the institutional intelligence of universities working as it should*.

The task is to recognise the politics, not to deny it.

Faculty resistance to AI literacy requirements often encodes a substantive concern about curricular coherence; the response is curricular dialogue, not bureaucratic mandate.

Faculty resistance to assessment redesign often encodes a labour concern; the response is workload reform, not exhortation.

Senate defence of inherited evaluation often encodes a substantive concern about academic standards; the response is to demonstrate that the new evaluation regime meets those standards more effectively under AI conditions.

Each veto point is a conversation, not an obstacle.

The reflexive condition of section 4.3 enters here as a discipline. A university whose senior leadership asks faculty to become orchestrators of AI-mediated learning while exempting itself from a corresponding redesign of its own labour and governance fails the reflexive test and **forfeits** the credibility it needs to do the political work. *The orchestration logic of* **Meira**, **Neves** and **Braga** (2025) *applies to leadership too*: senior leaders are themselves orchestrators of an organisational learning process, and their craft is the design of the organisational sequence, **not** the imposition of a managerial template. The recognition is uncomfortable for leaders who entered office under an older managerial model. It is also the precondition of any reform that **survives** the political process.

11.3 External Arenas

External actors carry their own interests. Vendors lobby for procurement rules that favour their architectures. Governments use AI governance to advance broader political agendas, including agendas not aligned with the public-mission framing of higher education. Employers seek alignment of credentials with their immediate demand, which may be narrower than the institution's educational mission. Donors press for visible innovation, which may not correspond to the institutional priorities the reform requires. Professional bodies defend their gates, sometimes for substantive reasons and sometimes from regulatory inertia. These interests are neither pathological nor uniform; they must be engaged, balanced and, where necessary, refused.

The international layer adds further complexity. The European AI Act sets a regulatory template whose effects will propagate through global markets via the so-called Brussels effect. The **NIST** (2023, 2024) framework provides a different template whose adoption is similarly extra-jurisdictional. National AI policies in China, India, Brazil and other major systems are developing on **different** trajectories. International credential recognition, research collaboration, student mobility and platform access are all sensitive to these regulatory differentials. Universities operating in multiple jurisdictions — whether through

campuses, partnerships or online provision — must navigate the differentials explicitly rather than assume a single global framework.

The **Boston Consulting Group** (2026) and Deloitte (**Clark** et al., 2026) analyses describe these external arenas from an industry perspective; **UPCEA** (2025) describes the adult- and professional-learning ecosystem; the Higher Education Policy Institute (**Stephenson** and **Armstrong**, 2026) describes the student-side dynamics. The signals are **not** unanimous. They are the input from which institutional strategy is constructed. The political work of higher education leadership is to **read** those signals, **situate** them in the institution's mission, **engage** the external actors as substantive interlocutors rather than as objects of communication, and **make explicit choices** about which external alignments to pursue and which to refuse.

11.4 Legitimate Negotiation Spaces

The institutional task is to create legitimate negotiation spaces. Faculty must be treated as co-designers, not implementation barriers. Students must be represented as rights-bearing participants, not merely users. Professional staff must be recognised as carriers of institutional knowledge rather than as compliance functions. Unions, where they exist, are partners in workload models, promotion reform and academic-work redesign. Transparent governance documents, independent learner and faculty voice, sunset clauses on rules whose value has expired, and public reporting of how AI-native reform changes outcomes are the building blocks of legitimacy.

The negotiation spaces are not symmetrical and should not pretend to be. Senior leadership has authority faculty do not have; faculty have professional autonomy senior leadership does not have; students have the moral standing of those whose education is at stake; professional staff have operational knowledge senior leadership cannot replicate; unions have legal standing the institution cannot ignore. Recognising the asymmetries explicitly is the condition of legitimate negotiation. *Pretending they do not exist generates the cynicism that erodes institutional culture*.

The architecture of negotiation should also be temporally extended. Reform that is settled in a single planning cycle and then closed encounters new conditions it cannot accommodate. The implementation horizons of section 9 are designed to keep the negotiation open across the cycle: ninety-day reviews of the prototypes, twelve-month reviews of the assessment redesign, three-year reviews of the workload models, ten-year reviews of the institutional architecture. The temporal extension is not a bureaucratic burden. It is the political infrastructure of credible reform.

11.5 Who Receives the Gains From Creative Destruction?

The deepest political question is who receives the gains from creative destruction inside the university. If AI removes routine work, does the saved time become mentoring, research, accessibility, public engagement and better feedback — or does it become budget reduction and adjunctification? If new orchestration roles emerge, do they accrue to a thin layer of platform-facing managers or to a broader community of educators, librarians and learning designers? *The Schumpeterian image of creative destruction is incomplete without a distributive question*; AI-native transformation will be judged, in part, by whom it makes more secure and whom it makes more precarious.

The distributive question is also a competence question. The orchestration roles that emerge in an AI-native institution require capabilities — pedagogical design, assessment expertise, data governance, professional development, learning research — that are presently concentrated in a relatively **small** population of professional staff and faculty. If those roles accrue **without** deliberate distribution, the institution will reproduce internally the stratification it is supposed to address externally. The countermove is a **deliberate development** of those capabilities across the institution, including in adjunct and contract roles, and a recognition of those capabilities in promotion, compensation and career architecture.

Table 5 maps the common political arenas of AI-native transformation and the questions they raise. The table is not a checklist. It is a **discipline** for thinking about which conversations the institution must hold open and what the substantive content of those conversations should be. Each arena, in each institution, will have its own history and its own actors. The table identifies the structural questions that any serious reform must address regardless of local detail.

Table 5. Political arenas in AI-native transformation

Arena	Likely conflict	Governance question
Faculty workload	Assessment redesign and mentoring require more time.	What work is reduced, recognised, resourced or redistributed?
Academic freedom	Platform-mediated teaching may constrain professional judgement.	How can faculty question models, vendors and institutional policies?
Student rights	AI creates new forms of monitoring, accusation and dependency.	What rights exist to guidance, access, privacy, appeal and portability?

Departmental autonomy	Cross-curricular AI literacy may cut across disciplines.	What is common to all graduates and what is domain-specific?
Procurement	Vendors offer fast solutions with hidden lock-in.	What audit, exit, interoperability and data protections are non-negotiable?
Accreditation	Existing standards may constrain new models.	How can evidence of learning be evaluated without metric reduction?
Employer influence	Workforce alignment may narrow educational purpose.	How are professional pathways balanced with civic and intellectual formation?
Public finance	AI may be used to justify cost cutting.	How are efficiency gains reinvested in human support and access?
Internal labour market	AI may destroy routine tasks and create orchestration roles unevenly across faculty, adjuncts and professional staff.	How are gains redistributed, new roles recognised, precarious work protected and career pathways preserved?

12. Global Diversity and the Global South as a Site of Theory

12.1 Beyond Deficit and Imitation

A global policy framework must avoid treating the Global South as a site of deficit or adoption. Low enrolment ratios, infrastructure constraints, linguistic diversity, limited fiscal capacity and precarious academic labour create severe challenges. They also make **visible** design questions that high-resource systems can ignore:

> **how** to deliver coherent formation at scale with thin human resources, **how** to honour multilingualism in a world of English-dominant models, **how** to recognise prior learning, **how** to combine formal and non-formal pathways, **how** to assemble public infrastructure under sovereignty constraints, and **how** to organise lifelong learning as a core function rather than a marginal revenue stream.

The treatment in this section deliberately exceeds the gestures common in policy documents because the Global South contains the majority of the world's present and

future students, and the framework cannot be considered serious if it addresses their conditions **only** as an afterthought.

The deficit framing is empirically inaccurate as well as politically problematic. Many Global South systems have built large public universities under fiscal and political constraints that high-income systems have **never** faced; the institutional intelligence accumulated in those systems is substantial. The imitation **framing** — that the trajectory of the Global South is to imitate, with delay, the trajectory of the high-income centre — is empirically falsified by mobile-first banking, mobile-first health systems, mobile-first agricultural extension and a number of other infrastructures in which the Global South **leads**. Higher education will likely follow the same pattern. The deliberate engagement of Global South challenges as theoretical priorities is therefore both substantive and politically necessary.

UNESCO's 2026 framing of higher education as a public mission provides a starting point but cannot be more than that. International organisations operate at scale and across jurisdictions; their templates are general by construction. The substantive design work has to be done **locally**. The role of an international framework is to provide the vocabulary, the comparative data, the regulatory architecture and the public-good orientation within which local design proceeds. The role of national and institutional actors is to **do** the design.

The triadic framework of **Meira** et al. (2026) and the phygital ontology of **Meira** and **Meira** (2026) are particularly relevant to Global South design because they were developed in part from Global South conditions. The work was produced in a Brazilian institutional context, with explicit attention to the asymmetries of the global digital economy and to the role of public mission under conditions of resource constraint. The texts are **not** Northern frameworks imported into the South; *they are South-originating conceptual instruments that the present paper extends to higher education globally*. Their theoretical priorities — the role of public infrastructure, multilingual capacity, recognition of prior learning, sovereign data governance — are **not** concessions to local conditions. They are general design principles whose local salience is highest in the South and whose long-term relevance is universal.

12.2 Specific Challenges: Infrastructure, Language, Labour, Sovereignty

Four pressures recur across Global South contexts and deserve more than a gesture.

> The **first** is **infrastructure**: intermittent connectivity, shared devices, energy constraints, limited cloud capacity and dependence on imported platforms. Mobile-first design is not a stylistic preference; it is the dominant access modality. The implication is that any educational architecture that assumes a stable broadband connection, a personal computer and a quiet home study environment is excluding

the majority of the global learner population. The vignettes in sections 7.6 and 7.7 illustrate concrete responses, but the implication is general: AI-native design must work in the conditions in which most students will actually study.

The **second** is **language and epistemic representation**. Models trained predominantly on English and a few dominant languages systematically marginalise Indigenous, regional and minority languages and the knowledge traditions they carry. Multilingual AI is a policy demand, not a product feature. The phygital triadic framework of **Meira** and **Meira** (2026) makes the issue analytically explicit: the Networked Digital dimension is grounded in computability and the Networked Social in communicability, and both are filtered through the corpora on which contemporary models are trained.

A language poorly represented in those corpora is structurally disadvantaged in the digital dimension regardless of its richness in the social and physical dimensions of phygital life. The policy task is therefore **not only** the localisation of interfaces; **it is** the active development of corpora, models, evaluation benchmarks and institutional contexts in which under-represented languages can be served by AI rather than disadvantaged by it.

The **third** is **academic labour**. In many Global South systems, large fractions of teaching are performed by adjunct or contract faculty under conditions that AI-mediated platforms **can** intensify rather than relieve. The digital-Taylorism risk analysed in **Meira**, **Neves** and **Braga** (2025) is acute in these systems because the labour conditions already trend toward **fragmentation**, **precarity** and **dashboard management**. *The political-economy implications are stark*. AI-native reform without explicit labour protections will entrench the precarity it claims to address. The political-economy framework of section 11 applies in the South with **greater urgency** and **lower margin for error** than in the high-income centre.

The **fourth** is **sovereignty**. Model providers, data centres, regulatory regimes and credential standards are predominantly located in the Global **North** (1990); without deliberate policy, AI-native transformation reproduces and deepens the colonial geometry of digital infrastructure. The platform analysis in **Meira** (2025a) identifies the structural mechanism: *when a global digital infrastructure orchestrates value across a value chain, the political-economy gains accrue to whichever jurisdiction hosts and regulates the orchestration layer*.

Sovereignty is therefore not a defensive concern; it is a **precondition of public capacity**. Public AI infrastructure, hosted under national or regional jurisdiction with auditability and public-interest terms, is the structural answer. Whether such

infrastructure will exist by 2035 is a political question that the international organisations alone cannot settle.

12.3 Design Principles for Low-Resource Contexts

Five design principles follow from the analysis above.

First, **public infrastructure**: national and regional consortia should procure, host or develop core AI capabilities under public-interest terms. The instrument may be a national model, a regional federation, a hybrid arrangement with private providers under audit and exit conditions, or a non-profit infrastructure governed by educational consortia. *The form matters less than the principle: the educational AI infrastructure must not be entirely outsourced to private platforms whose interests are not aligned with the public mission*.

Second, **multilingual by default**. Any platform used at scale must support local languages with quality benchmarks specified by local academic communities. The benchmark specification is the substantive instrument. Without local academic standards for what constitutes adequate model performance in a given language, the linguistic question **reduces to** translation quality, which is necessary but insufficient. Adequate performance requires translation, idiomatic accuracy, register sensitivity, terminological discipline and the capacity to evaluate domain-specific outputs in the language of the discipline. None of this is automatic; all of it requires sustained academic engagement.

Third, **mobile and low-bandwidth design**: assessment, mentoring, content and credentials must work in conditions that are not those of the elite campus. The implication is architectural. Educational systems designed assuming desktop access, stable connectivity and abundant device time must be redesigned for the conditions in which most learners actually study. The vignettes in sections 7.6 and 7.7 are illustrative; *the structural point is that the elite-campus architecture is not portable to mass-access conditions without distortion*.

Fourth, **decent academic work**: labour rights, workload models and professional development for educators are part of AI policy, not a separate file. The labour conditions of teachers and professional staff are the operational substrate on which any pedagogical reform must run. *A reform that does not address those conditions is a reform that will fail at the level of execution regardless of how well it is designed at the level of intent*.

Fifth, **recognition of prior and informal learning**: lifelong learning architectures must credit what learners already know, including knowledge formed outside formal

institutions. The **Council of the European Union**'s (2022) Recommendation on micro-credentials provides one instrument; comparable instruments are emerging in other systems. The principle is that the credential must serve the learner, not the bureaucratic convenience of the certifying institution.

12.4 The South as Theoretical Priority

Treated seriously, the Global South is **not a** peripheral case **but a** theoretical priority. The challenges of multilingualism, public infrastructure, mobile-first pedagogy, low-cost quality, recognition of prior learning and precarious labour are not local difficulties to be addressed once the general framework is settled. *They are the conditions under which higher education will be delivered for the majority of the global student population over the period 2026–2035*. A framework that is robust to those conditions is **general**. A framework that requires elite-campus conditions to function is **parochial**.

The implication is methodological. The development of theory about AI-native higher education should be substantively engaged with Global South cases and Global South theorists, not as supplementary illustration but as primary input. The triadic intelligence framework of **Meira** et al. (2026) and the phygital ontology of **Meira** and **Meira** (2026) are partial instances of this methodological commitment. Their further development requires engagement with parallel work originating in other Global South contexts — work that is presently underrepresented in the international literature for reasons that are themselves structural.

International organisations have a role, but their guidance must support local capacity rather than export generic templates. **UNESCO**'s 2026 framing is a starting point. The **OECD**'s Digital Education Outlook (**OECD**, 2026) provides comparative analysis whose Northern bias must be read against the institutional conditions in which it was produced. The **Boston Consulting Group**'s analysis (2026) similarly carries the methodological constraints of its source. None of this **disqualifies** the work; all of it requires that local actors read the international literature with explicit attention to whose conditions and whose interests the literature reflects. The substantive South-South engagement that has emerged across health systems, agriculture and mobile finance in the past two decades has not yet been replicated at the same scale in higher education. The AI moment is an opportunity to begin.

13. Risks, Refusal Conditions and Governance Responses

13.1 Seven Major Risks Read Through the Red-Swan Register

AI-native transformation creates real **opportunities** and serious **risks**. A credible framework must specify both what should be built and what should be prevented, constrained or refused. **Seven** risks recur. They are listed below with brief mechanism and response; Table 6 specifies the architecture in full. The risks are read through the red-swan register of **Meira** (2025b) because their character is not merely high-probability adverse outcomes amenable to standard risk management. Several of them have the structure of **ontological displacement**: they change the categories within which the institution operates, not merely the outcomes that the institution must manage.

Platform capture: dependence on a small number of vendors for models, analytics, identity systems and credential infrastructures.

> The platform analysis in **Meira** (2025a) identifies platform capture as the structural endpoint of value-chain disintermediation absent deliberate countermeasures.
>
> **Governance response**: open standards, audit rights, public procurement rules, exit clauses and federated alternatives.

Stratification: intensive human-AI education for elites and thin automation for mass learners.

> The mechanism is documented by **Stephenson** and **Armstrong** (2026) in the differential expectations of students at well-resourced and under-resourced institutions.
>
> **Governance response**: equity metrics, public investment, human-support requirements and quality assurance.

Cognitive outsourcing: learners delegate reasoning before developing judgement, undermining the very expertise the system claims to produce.

> The **OECD**'s Digital Education Outlook (**OECD**, 2026) documents the risk in detail; it is also the central concern of the **Bearman** et al. (2024) analysis of evaluative judgement.
>
> **Governance response**: AI literacy, deliberate AI-free tasks, process assessment and pedagogies of verification.

Assessment collapse: unsupervised outputs no longer prove individual competence.

The **Boud** and **Falchikov** (2006) alignment principle implies that the credential will eventually lose its signal value.

Governance response: oral defence, authentic tasks, simulations, supervised work and portfolios.

Faculty degradation and digital Taylorism: AI used to deskill, monitor, score or replace educators, and to fragment academic judgement into dashboards.

The analytic vocabulary for this risk is developed in **Meira**, **Neves** and **Braga** (2025).

Governance response: role redesign, workload protection, staff development, academic-freedom safeguards, participatory governance.

Regulatory capture: incumbents or vendors shape rules to protect narrow interests. ‘

The history of outcomes-based assurance documented in section 10.3 is the standard case.

Governance response: independent review, learner representation, public consultation and sunset clauses.

Mission drift: AI strategy framed mainly as efficiency or workforce supply, eroding the civic, scientific and cultural purposes of higher education that **Marginson** and **Barnett** (2011) analyse as the substantive content of the institution.

Governance response: public value tests, civic indicators, research protection and mission-based funding.

Table 6. Major risks and governance responses

Risk	Mechanism	Governance response
Platform capture	Dependence on proprietary models, analytics and credential systems.	Open standards, audit rights, public procurement rules, exit clauses and federated alternatives.
Stratification	High-quality human-AI education for elites and thin automation for mass learners.	Equity metrics, public investment, human-support requirements and quality assurance.
Cognitive outsourcing	Learners delegate reasoning before developing judgement.	AI literacy, deliberate AI-free tasks, process assessment and pedagogies of verification.

Assessment collapse	Unsupervised outputs no longer prove individual competence.	Oral defence, authentic tasks, simulations, supervised work and portfolios.
Faculty degradation and digital Taylorism	AI used to deskill, monitor, score or replace educators and to fragment academic judgement into dashboards.	Role redesign, workload protection, staff development, academic-freedom safeguards, participatory governance.
Regulatory capture	Incumbents or vendors shape rules to protect narrow interests.	Independent review, learner representation, public consultation and sunset clauses.
Mission drift	AI strategy framed mainly as efficiency or workforce supply.	Public value tests, civic indicators, research protection and mission-based funding.

13.2 Refusal as Governance Capacity

AI-native governance must include refusal as a capacity. Institutions should be able to say no to uses of AI even when they are technically feasible and economically attractive.

> Pervasive surveillance, affective manipulation, automated disciplinary decisions, opaque admission scoring, unaccountable behaviour prediction and the wholesale outsourcing of pedagogical judgement to vendor algorithms are candidates for refusal.
>
> **Refusal is not technophobia**. It is an act of institutional self-definition that distinguishes universities from the platforms that imitate them.

The category of red-swan disruption analysed in **Meira** (2025b) is relevant here. Some technologies do not merely change the rules of the game; *they change which institutions can credibly play it*. Refusal is the institutional posture that prevents universities from being silently demoted to the residue of platform markets. The refusal is exercised at the level of procurement, governance, public communication and, where necessary, regulatory advocacy. It is not a single decision but a sustained institutional disposition that the political-economy framework of section 11 specifies in operational form.

Refusal also has a positive form: the explicit articulation of what the institution is for. **Marginson**'s (2016) common-good framing, **Barnett**'s (2011) defence of *being a university*, **Kerr**'s (1963) account of the multiversity as a precarious truce among incompatible

communities and the long tradition of educational philosophy from **Newman** (1852) onwards each provide vocabularies for the positive articulation. The institutional task is to make those vocabularies **operational** under contemporary conditions, so that refusal is not only the rejection of specific platform uses but the active enactment of an alternative institutional model.

13.3 Independent Oversight and Democratic Accountability

Refusal needs structure. Independent review boards, learner and faculty voice, public consultation, sunset clauses on rules whose value has expired, and statutory protection of academic freedom under AI-mediated conditions are the institutional architecture of refusal. Without them, refusal becomes **either** rhetorical or arbitrary. With them, it becomes a credible part of how universities are governed. The architecture should be designed **against** the regulatory capture pattern documented in section 10.3: **independent** in fact, not only in name; **consultative** in substance, not only in form; **sunsetted** in such a way that ossified rules can be dismantled without political crisis.

The democratic-accountability dimension is consequential. Universities are public institutions in most jurisdictions and quasi-public in others. Their accountability to the publics they serve is not a residual concern but a constitutive one. AI-mediated decisions about access, evaluation, progression and certification touch the public interest at a scale and depth that conventional administrative law has not yet caught up with.

> The institutional architecture of refusal and oversight must be designed against the standards that will eventually be applied to public-interest AI governance more broadly, including those emerging through the EU AI Act and comparable frameworks.

The international layer adds a final note. Refusal in one institution under one regulatory regime is operationally limited if the global infrastructure on which the institution depends operates under different terms. The Global South discussion in section 12 is directly relevant: sovereign capacity to refuse is not equally distributed, and the political-economy work of building such capacity is part of what AI-native transformation requires.

> The architecture of refusal is therefore **not an** internal university concern alone; **it is** a multi-level governance challenge whose operational form depends on the alignment between institutional, national and international layers.

14. Scenarios for 2035

14.1 Adaptive Public Renewal

Public universities, regulators and governments **invest** in open AI infrastructure, faculty development, modular credentials, authentic assessment and lifelong learning. Universities **retain** legitimacy because they demonstrate learning quality, equity gains, civic contribution and accountable governance. AI is **integrated** as cognitive infrastructure under public-interest terms; private platforms **participate** but do not dominate.

> The scenario corresponds to the upside feedback loop of section 8.3 and the public-mission framing of **UNESCO** and **Marginson** (2016). It is not the most probable scenario absent deliberate effort, but it is a coherent possibility that the institutional, organisational and ecosystem layers can be aligned toward.

The scenario has identifiable preconditions. It requires sustained public investment under fiscal conditions that have to be politically defended; coherent regulatory architecture in the lineage of the EU AI Act and the **NIST** (2023, 2024) framework; faculty co-design rather than managerial imposition; learner voice as a structural feature of governance; and refusal as a recognised institutional capacity. Each precondition is contestable in its own political register. None of them are spontaneous outcomes of market dynamics. The scenario is a design proposal, in the sense of section 2.2, not a forecast.

14.2 Platform-Captured Massification

Higher education **expands** through low-cost AI-mediated platforms controlled by a small number of vendors. Access **increases**. Quality becomes **uneven**, learner data are **commodified**, human support is **thin**, credentials **proliferate**, trust **fragments**. Public institutions become regulators of, or appendages to, vendor platforms. The scenario corresponds to the **downside** feedback loop of section 8.3 and to the structural endpoint identified by the platform analysis in **Meira** (2025a) under conditions of value-chain disintermediation without public countermeasures.

> The scenario's plausibility derives from its consistency with present incentive structures. Vendors have capital, technical capacity and organisational coherence; public institutions have authority but limited capacity and contested political legitimacy; students and employers have incentives to use the platforms regardless of long-term institutional consequences.

The **trajectory** toward this scenario does not require any single actor to make a catastrophic decision; it requires only that the present configuration **continue**. Avoiding the scenario therefore requires explicit, sustained countermeasures rather than confidence in default outcomes.

14.3 Defensive Institutionalism

Universities restrict AI, preserve conventional assessment, defend degree monopolies and resist modularisation. Some institutions maintain excellence; many lose relevance as employers, learners and alternative providers build parallel credential systems.

The conservative reflex protects short-term legitimacy at the cost of long-term capacity. The scenario corresponds to the institutional response that **Meira** (2025b) identifies as characteristic of incumbents facing paradigmatic transcendence: a **defence** of existing categories that progressively **loses** traction as the categories themselves cease to organise the surrounding environment.

The scenario has historical precedent. Newspapers, broadcast television, traditional retail and several other incumbent sectors have followed analogous trajectories under previous waves of digital transformation. The pattern is well documented and the analytical literature is consistent: defensive institutionalism is a coherent strategy for elite institutions whose positional good is preserved by the institutional defence, and a high-risk strategy for non-elite institutions whose value proposition depends on continued ecosystem trust. The differential outcomes across the institutional spectrum are a likely consequence of this scenario rather than a single uniform result.

14.4 Mission-Oriented AI-Native Transformation

Institutions **reorganise** around public missions: climate adaptation, health systems, education, democratic resilience, industrial transformation, public administration, cultural production and regional development. AI **supports** research, teaching and partnership in service of mission. Credentials are **tied** to demonstrated capability. The scenario is **consistent** with the Mazzucato-influenced mission-economy literature and with the **UNESCO** (2026) public-mission *framing*. It is also consistent with the triadic-orchestration analysis of **Meira** et al. (2026), in which the **integrative** capacity to compose individual, social and artificial intelligence around concrete public problems becomes the substantive content of higher-education work.

The scenario implies organisational **reorganisation** away from the disciplinary department as the primary unit of institutional life and toward mission-oriented programmes, centres and partnerships. The disciplinary department **does not**

disappear; it becomes a horizontal resource on which mission-oriented work draws. Faculty workload, promotion, governance and resource allocation are **redesigned** around the new institutional architecture. The scenario is operationally demanding and politically difficult. It is also one of the few that addresses the legitimacy problem of section 3.2 in substantive form.

14.5 Fragmented Prestige Dualism

A small number of high-status universities **combine** AI with intensive human formation and preserve strong labour-market signals. A large number of institutions **adopt** AI tools but cannot **redesign** assessment, data governance or faculty roles. Learners experience radically **different** educational architectures depending on access. Equity policy becomes the central political question of higher education. The scenario is the most **consistent** with present trajectories under conditions in which the political economy of section 11 is **not** actively addressed: the elite institutions **implement** AI-native architectures with the resources required; the mass-access institutions **adopt** the tools without the architectural redesign because they lack the resources, the governance capacity or the political mandate to redesign.

The scenario also reproduces, at a new scale, the **stratification** pattern **Slaughter** and **Rhoades** (2004) document under academic capitalism: educational architecture sorts into prestige and access tiers; the prestige tier captures the qualitative gains of the new architecture; the access tier captures the cost reductions without the qualitative gains. The political consequences are familiar: educational stratification reinforces broader social stratification, with implications for democratic legitimacy, social mobility and the political support for public investment in higher education. The scenario is **not** stable. It produces the political conditions under which **one** of the other scenarios — public renewal, platform capture or defensive retrenchment — will eventually **predominate**.

14.6 Reading Combined Trajectories

The scenarios are not mutually exclusive. Different segments of higher education may move toward different configurations. Elite global universities may become mission-oriented and platform-enhanced. Private providers may move toward platform capture. Public systems may divide between adaptive renewal and defensive institutionalism. Equity-driven systems may consciously construct lifelong-learning architectures. The strategic question for any institution is not which scenario will prevail globally but which configuration its mission, demographics, regulatory environment and capacity can support without abandoning public value.

The scenarios should also be read together with the maturity model of section 9.6.

An institution at Level 1 (Enablement) is structurally exposed to the platform-captured massification scenario unless it accelerates to Level 2 (Alignment) and beyond.

An institution at Level 3 (Redesign) has the architectural foundation for either public renewal or mission-oriented transformation depending on the alignment of its ecosystem and institutional layers.

An institution at Level 5 (Public renewal) or Level 6 (Reflexive renewal) is constructing the conditions under which the optimistic scenarios are possible at all.

The combination of maturity model and scenarios is a strategic instrument, not a forecast: it specifies the decision space within which the institution can act.

The temporal register matters. **Meira** (2025b) argues that strategic planning under contemporary conditions cannot rely on forecasting; it must build option value, institutional reflexivity and the cognitive infrastructure required to recognise paradigmatic shifts when they occur rather than after they are complete.

The scenarios are tools for that work. They specify configurations the institution can recognise in itself and in its environment, and the implementation architecture of section 9 is designed to be compatible with several of them simultaneously rather than to commit to one.

15. A Research Agenda and Comparative Indicators

15.1 Ten Comparative Priorities

The concept of AI-native higher education should not become a slogan. It should become a comparative object of study. The research agenda proposed here is organised around ten priorities, each of which translates a substantive claim of the framework into an empirically tractable question. The priorities are not exhaustive; they are the points at which the framework most needs disciplined empirical scrutiny. They are also the points at which the credible deployment of the framework in institutional decision-making most depends on evidence the present literature does not yet provide.

First, **assessment validity under AI conditions**: what proportion of high-stakes assessment still relies on products that can be generated by AI without detection, what proportion uses authentic or supervised evidence, and how students are required to justify and defend their work. The methodological framework is provided by the work of Bearman and colleagues and by the older alignment-with-long-term-learning literature of **Boud** and **Falchikov** (2006); the empirical work has barely begun at scale.

Second, **AI literacy and meta-cognitive orchestration across the curriculum**: can graduates in law, education, medicine, design, administration and the humanities orchestrate AI use under their disciplines' standards? The methodological framework is provided by the triadic intelligence analysis of **Meira** et al. (2026) and the meta-cognitive orchestration framework of **Meira**, **Neves** and **Braga** (2025). The empirical question is whether the institutional architecture actually produces the orchestration capacity it claims.

Third, **human-support intensity**: where is human time allocated, and does automation increase access to meaningful support? The question is operationally consequential because the principal pedagogical claim of AI-native higher education — that AI is most valuable when embedded in sequences that include human mentoring, supervision and assessment — depends on the human-time budget actually being reallocated rather than reduced.

Fourth, **learner data sovereignty**: can students access, export and control their learning records, credential evidence and AI interaction data? The methodological framework is provided by the EU AI Act and comparable data-protection regimes; the empirical question is whether the rights are operationally exercisable.

Fifth, **platform dependency**: how dependent are institutions on proprietary models, opaque scoring systems and non-portable credentials, and what exit options exist? The platform analysis in **Meira** (2025a) provides the analytical framework; the empirical work requires comparative audit.

Sixth, **lifelong-learning integration**: is recurrent formation central to mission and finance? The **UPCEA** (2025) reporting describes the demand side; the supply-side empirical work is uneven.

Seventh, **public-mission alignment**: does AI transformation strengthen teacher education, public health, climate adaptation, democratic resilience, regional innovation and cultural work? The **Marginson** and **UNESCO** (2026) framing specifies the substantive content; the empirical work requires institutional-level measurement.

Eighth, **the political economy of prestige**: how do verified learning evidence, responsible data governance and equity outcomes reshape the geometry of institutional reputation? The **Slaughter** and **Rhoades** (2004) framework specifies the structural mechanism; the empirical work requires longitudinal analysis.

Ninth, **academic-labour transformation**: does AI reduce workload or intensify it, improve teaching quality or standardise it, create new professional categories or precarise existing ones? The methodological framework is provided by **Meira**, **Neves**

and **Braga** (2025); the empirical work requires labour-conditions surveys at institutional scale.

Tenth, **the governance of modular credentials**: which credentials are recognised, which can be accumulated, and how is quality assured? The **Council of the European Union**'s (2022) Recommendation provides one regulatory framework; comparative analysis across systems is required.

Table 7. Comparative indicators for AI-native higher education

Indicator	Core question	Evidence sources
Assessment validity	Can the institution credibly certify competence under AI conditions?	Assessment maps, audit samples, oral defence records, external moderation.
AI literacy and orchestration across curriculum	Is AI competence distributed beyond technical programmes?	Programme outcomes, syllabi, student portfolios, oral defences, process logs, graduate surveys.
Human-support intensity	Does AI increase or reduce meaningful human support?	Advising ratios, feedback data, mentor access, student experience evidence.
Learner data sovereignty	Can learners control, export and contest data and decisions?	Data policies, vendor contracts, portability mechanisms, appeal records.
Platform dependency	Can the institution audit, switch, interoperate and preserve records?	Procurement clauses, architecture maps, incident reports, exit tests.
Lifelong-learning integration	Is recurrent learning central to mission and finance?	Adult enrolments, modular pathways, recognition of prior learning, stackability.
Public-mission alignment	Does AI transformation serve collective problems?	Mission indicators, regional partnerships, civic projects, equity outcomes.
Prestige–learning alignment	Does reputation correspond to demonstrated learning quality?	External assessment, graduate performance evidence, employer and civic trust.

Academic-labour transformation	Are roles, workload and autonomy improved or degraded?	Workload studies, contracts, promotion criteria, staff surveys.
Credential coherence and portability	Are modular credentials trusted, assessed and transferable?	Credit transfer data, employer recognition, quality assurance, learner records.
Reflexive work redesign	Does the institution transform its own academic and professional work consistently with the capability it certifies?	Workload audits, promotion criteria, staff-development records, new role descriptions, collective agreements.

15.2 Indicators and Their Epistemic Limits

Indicators are useful and dangerous. Used as diagnostic signals supplemented by qualitative judgement, they support strategic conversation. Used as compliance targets, they generate the gaming and capture this paper warns against in section 10.3.

> The composition principle for research is therefore that **indicators must travel together**: assessment validity without academic-work redesign is suspect; data sovereignty without procurement reform is paper-thin; human-support intensity without equity disaggregation hides stratification.
>
> The strongest empirical work on AI-native transformation will be comparative, mixed-method and reflexive about its own metrics.

The reflexive condition applies here too. The research community that studies AI-native higher education is itself part of the institutional ecology it studies. Its institutional position, funding sources, methodological commitments and disciplinary identities shape what it can and cannot see.

> Disciplined transparency about those shaping conditions — the methodological equivalent of the data-rights statement institutions owe their learners — is part of what credible empirical work on AI-native higher education requires.

The danger is that the research literature on AI-native transformation **will replicate** the platform-capture dynamic at the level of analytical attention: *a few well-resourced research centres define what gets studied, how it gets measured and what counts as evidence, with predictable consequences for what the institutional decision-makers learn to attend to*.

16. Conclusion: Designing the University That Deserves to Exist

Higher education will not disappear after generative AI.

> *The more serious possibility is that significant parts of it will become less credible, less equitable and less public if transformation is left either to institutional inertia, to digital Taylorism, or to platform markets alone*.

The university that deserves to exist **after** generative AI is one that can **articulate** what its graduates can do, **prove** how they were formed, **govern** the data of their learning, **redesign** the work of its educators, **recognise** prior and lifelong learning, **refuse** uses of AI that violate its mission, and **contribute** to the public problems of its century.

This paper has argued that higher education should be understood as a sociotechnical hyperobject, with the hyperobject metaphor used cautiously as a discipline against single-surface reform rather than as an ontological claim.

It has proposed the AI-enabled / AI-native distinction, with nine constitutive features of the AI-native institution.

It has grounded the framework in the triadic intelligence ontology developed in **Meira** et al. (2026) and in the meta-cognitive orchestration analysis of **Meira**, **Neves** and **Braga** (2025).

It has offered eight vignettes that translate the framework into concrete daily life across disciplines, modalities and contexts.

It has proposed a four-horizon implementation architecture and a seven-level maturity model that answer the Monday-morning question.

It has acknowledged the dual risk of outcomes-based assurance and the structural pull toward compliance theatre.

It has mapped the political economy of change inside universities, the differentiated trajectories of the Global South, and the seven major risks that justify refusal as a governance capacity in the red-swan register of **Meira** (2025b).

It has sketched five scenarios for 2035 and a ten-priority research agenda with comparative indicators.

Underneath the surface of this **architecture**, the **substrate** of **Meira** and **Meira** (2026) is doing analytical work. The phygital triadic ontology — Physical, Networked Digital, Networked Social — and the informaticity substrate that grounds it — computability (**Turing**, 1936), communicability (**Shannon**, 1948), controlability (**Wiener**, 1948) — provide the conceptual base on which the higher-education-specific argument is built.

The platform-orchestration analysis in **Meira** (2025a) provides the political economy of how value circulates and concentrates in environments where general-purpose digital infrastructure orchestrates previously bundled functions.

Together, the five framework texts referenced throughout this paper constitute the conceptual scaffolding that connects the educational claim to the broader analysis of cognitive work, institutional design and strategic rupture under contemporary conditions.

The reflexive test remains decisive. A university cannot credibly prepare students for meta-cognitive orchestration while organising its own labour as programmed compliance, invisible overload or dashboard-mediated control.

The creative destruction required by generative AI must occur within the university, not only around it. The transition from programmed labour to meta-cognitive orchestration is the substantive content of that destruction, and the institutional architecture that delivers it must be designed against the digital-Taylorism trajectory rather than as a tacit acceleration of it. The political economy of section 11, the policy architecture of section 10, the implementation horizons of section 9 and the refusal capacity of section 13 are the operational instruments through which that institutional design is realised.

The stakes are high because AI changes not only the means of education but the social contract around it. If machines can produce fluent text, solve routine problems, simulate explanation and support personalised practice at scale, universities must clarify what their certificates mean, what their teachers do, what their students learn, what their research achieves, what their data protect and what their institutions refuse. The Schumpeterian metaphor of creative destruction captures the scale of the moment but **underspecifies** its distributive content; whether the destruction is borne by precarious labour and disadvantaged learners while the creation is captured by elite institutions and global platforms is **not** a market outcome **but** a political and institutional choice.

Higher education has redesigned itself before.

> **Trow**'s (1973) expansion analysis, **Kerr**'s (1963) multiversity architecture, **Clark**'s (1983) comparative institutional theory, **Barnett**'s (2011) philosophical defence of *being a university*, **Marginson**'s (2016) recovery of the common good, **Slaughter** and **Rhoades**'s (2004) diagnosis of academic capitalism, and the long policy literature on quality assurance, credential recognition and academic labour each document an episode in which the institution remade itself under conditions that earlier generations found similarly disorienting.

The present moment is not unique in its character; it is unique only in its specific content. The institutional intelligence required to navigate it is continuous with the institutional intelligence required to navigate previous moments. *What is new is the speed of the*

*substrate change, the scale of the platform pressure and the simultaneity of the demographic, fiscal, technological and political pressures with which the institution is asked to conten*d.

The university after generative AI should not be a nostalgic campus defended against machines, nor a platform that automates learning at scale while calling itself education.

> *It should be a human institution redesigned for a world in which cognition is technologically mediated*.
>
> Its purpose is not to compete with the machines, but to make sure that the humans who use them remain capable of judgement, accountability and public reason.

The work is not done.

The framework offered here is a **starting point** for the design conversation that institutions, regulators, students, faculty and publics now have to hold. The conclusion of any policy review that takes its own subject seriously is the same: **read the text, argue with it, refuse what should be refused, and build the institution the world needs**.

References

Acemoglu, Daron, and Pascual Restrepo. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives*, vol. 33, no. 2, 2019.

Autor, David H., and Neil Thompson. *Expertise*. National Bureau of Economic Research, Working Paper no. 33941, 2025.

Barnett, Ronald. *Being a University*. Routledge, 2011.

Bearman, Margaret, Joanna Tai, Phillip Dawson, David Boud, and Rola Ajjawi. “Developing Evaluative Judgement for a Time of Generative Artificial Intelligence.” *Assessment & Evaluation in Higher Education*, vol. 49, no. 6, 2024.

Boston Consulting Group. *How AI Can Help Higher Education Capture a Once-in-a-Generation Opportunity*. 2026.

Boud, David, and Nancy Falchikov. “Aligning Assessment with Long-Term Learning.” *Assessment & Evaluation in Higher Education*, vol. 31, no. 4, 2006.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*. Stanford Digital Economy Lab, 2025.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. “Generative AI at Work.” *The Quarterly Journal of Economics*, vol. 140, no. 2, 2025.

Clark, Burton R. *The Higher Education System: Academic Organization in Cross-National Perspective*. University of California Press, 1983.

Clark, Cole, Megan Cluver, Tiffany Fishman, and Hayley Glatter. *2026 Higher Education Trends*. Deloitte Insights, 2026.

Council of the European Union. “Council Recommendation of 16 June 2022 on a European Approach to Micro-credentials for Lifelong Learning and Employability (2022/C 243/02).” *Official Journal of the European Union*, C 243, 2022.

Drucker, Peter F. *Landmarks of Tomorrow*. Harper, 1959.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. “GPTs Are GPTs: Labor Market Impact Potential of LLMs.” *Science*, vol. 384, no. 6702, 2024.

European Parliament and Council. “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).” *Official Journal of the European Union*, L 2024/1689, 2024.

Fuggetta, Alfonso. "3+ 1 Challenges for the future of universities." *Journal of Systems and Software* 85.10 (2012): 2417-2424.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller. *Artificial Intelligence and the Labor Market*. National Bureau of Economic Research, Working Paper no. 33509, 2025.

Kerr, Clark. *The Uses of the University*. Harvard University Press, 1963.

Marginson, Simon. *Higher Education and the Common Good*. Melbourne University Publishing, 2016.

Massenkoff, Maxim, and Peter McCrory. *Labor Market Impacts of AI: A New Measure and Early Evidence*. Anthropic Economic Research, 2026.

Meira, Silvio, et al. *As Três Inteligências, Revisitadas: Tríade Sinérgica, Redes de Agentes e a Orquestração Metacognitiva do Trabalho na Era Pós-LLM*. Porto Digital, 2026. To be published.

Meira, Silvio, *Everything Platforms*. SSRN, http://dx.doi.org/10.2139/ssrn.5160171, 2025a.

Meira, Silvio. *Red Swans: Ontologies of the Unthinkable in the Age of Strategic Collapse*. SSRN , http://dx.doi.org/10.2139/ssrn.5486446, 2025b.

Meira, Silvio, André Neves, and Carlos A. Primo Braga. *From Programmed Labor to Meta-Cognitive Orchestration*. SSRN, http://dx.doi.org/10.2139/ssrn.5329936, 2025.

Meira, Silvio, and Luciano Meira. *Phygital Space: Ontology, Informaticity, and the Architecture of Emergent Reality*. TDS.company / cesar.school / PROZ, 2026. To be published.

Morton, Timothy. *Hyperobjects: Philosophy and Ecology after the End of the World*. University of Minnesota Press, 2013.

National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1, U.S. Department of Commerce, 2023.

National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. NIST AI 600-1, U.S. Department of Commerce, 2024.

Newman, John Henry. *The Idea of a University*. 1852.

North, Douglass C. *Institutions, Institutional Change and Economic Performance*. Cambridge University Press, 1990.

OECD. *OECD Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education*. OECD Publishing, 2026.

Polanyi, Michael. *The Tacit Dimension*. Routledge & Kegan Paul, 1966.

Schumpeter, Joseph A. *Capitalism, Socialism and Democracy*. Harper & Brothers, 1942.

Selwyn, Neil. *Should Robots Replace Teachers? AI and the Future of Education*. Polity Press, 2019.

Shannon, Claude E. "A Mathematical Theory of Communication." *Bell System Technical Journal*, vol. 27, 1948.

Slaughter, Sheila, and Gary Rhoades. *Academic Capitalism and the New Economy: Markets, State, and Higher Education*. Johns Hopkins University Press, 2004.

Stephenson, Rose, and Charlotte Armstrong. *Student Generative Artificial Intelligence Survey 2026*. HEPI Report 199, Higher Education Policy Institute, 2026.

Trow, Martin. *Problems in the Transition from Elite to Mass Higher Education*. Carnegie Commission on Higher Education, 1973.

Turing, Alan M. "On Computable Numbers, with an Application to the Entscheidungsproblem." *Proceedings of the London Mathematical Society*, 2nd ser., vol. 42, 1936.

Turing, Alan M. "Computing Machinery and Intelligence." *Mind*, vol. 59, no. 236, 1950.

UNESCO. *Transforming Higher Education: Global Collaboration on Visioning and Action*. UNESCO, 2026.

UPCEA. *Predictions 2026: Insights for Online and Professional Education*. University Professional and Continuing Education Association, 2025.

Wiener, Norbert. *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT Press, 1948.


tds.company