SEMPER VERUM PARTNERS

THE RIGHT ANSWER, AND WHO CAN AFFORD IT

The Cost of Being Sure

Higher education’s AI problem is an assessment problem, and nobody has priced it

Jack Jennings · Vince Wilczynski

A Semper Verum White Paper · October 2026

Motivation for This White Paper

Semper Verum Partners works with institutional leaders operating where academic mission, demographic change, financial sustainability, technological disruption, and public trust intersect.

Our work focuses on questions that require business realities and academic realities to be reasoned about together. We are particularly interested in challenges that sit between organizational strategy and educational practice: the places where governance decisions, financial constraints, instructional design, and institutional purpose meet.

Artificial intelligence is one such challenge.

Most discussion of AI in higher education has focused on technology, pedagogy, ethics, or workforce preparation. These are important questions. But they often leave unexamined the economic and institutional assumptions on which higher education itself rests. This paper examines one of those assumptions. A degree has value because employers believe it contains information about what its holder can actually do. The question addressed here is simple: **If AI makes the appearance of competent work easy to produce, what does it cost to preserve confidence that competence is real?**

Executive Summary

Artificial intelligence has exposed a problem higher education did not create but can no longer avoid. The issue is not whether students use AI. The issue is whether institutions can still verify what their degrees certify.

A credential derives its value from the information it carries. Employers pay a wage premium for a degree because they believe it contains information about what its holder can actually do. That value depends on verification: an institution's ability to credibly attest that a student can perform at the level the credential implies. Where assessment occurs under observation, such as laboratories, clinical rotations, studio critiques, oral examinations, or proctored assessments, verification remains intact. Where assessment occurs out of sight, AI has weakened the connection between completed work and demonstrated competence.

This is why higher education's AI problem is fundamentally an assessment problem rather than an integrity problem. Academic integrity asks who cheated. Verification asks what an institution can still attest to. And the exposure is not distributed by discipline: it follows the assessment design. Programs built around supervised demonstration remain comparatively resilient; programs dependent on work produced outside observation face the largest verification challenge.

Most institutions have recognized one new cost: providing access to AI tools. Far fewer have recognized the second: restoring verification through redesigned assessment, supervised demonstration, oral examination, defended portfolios, and additional instructional capacity. Those costs are not equally fundable. Technology spending is visible, vendor-based, and measurable. Verification spending is instructional, labor-intensive, and recurring.

As a result, the principal institutional consequence of AI may not be adoption but divergence. The institutions most in need of restored verification may be the least able to afford it.

Meanwhile, the pedagogical answer is increasingly clear. Major inquiries at MIT, Cornell, Stanford, Chicago, and Duke are converging on a similar principle: students must learn to think both with AI and without it. The unresolved question is not educational. It is financial: **what does it cost to know that learning occurred, and who can pay?**

Key Findings

1. Higher education's AI challenge is fundamentally a verification problem. An institution can maintain strong integrity policies while losing confidence in what its credentials attest to.

2. The exposure follows assessment design, not discipline. The key distinction is not humanities versus engineering. It is whether graded work is produced under observation.

3. Most institutions have priced AI access, but not verification restoration. The sector increasingly budgets for AI licenses and infrastructure. Far fewer institutions have budgeted for oral examinations, smaller sections, portfolio defenses, and supervised practice.

4. Verification is primarily an instructional expense. Technology access can be purchased. Verification must generally be staffed. Most forms of credible verification scale with contact hours and instructional labor.

5. The likely outcome is institutional divergence. The pedagogical answer is increasingly known. The operational challenge is affordability, and the institutions most in need of restored verification may be least able to afford it.

Questions for Institutional Leaders

> **Verification.** What percentage of required coursework is completed under observation? Where is verification strongest, and where is it weakest?
>
> **Assessment.** Which programs remain highly dependent on work produced out of sight? What evidence exists that graduates can perform unaided at the level their credentials imply?
>
> **Economics.** Does the institution's financial model include assessment-restoration costs, or only AI-access costs? If instructional spending remains flat, how is verification being restored?
>
> **Market signals.** What would employer confidence erosion look like before enrollment declines appear? What independent evidence do employers have that graduates possess the capabilities the institution certifies?

1. Introduction

A degree certifies that someone can do something. Where that certification rests on supervised performance, such as a lab bench, a clinic, a studio, a practicum, or a proctored examination, it still holds. Where it rests on work produced out of sight, it is eroding. Restoring it is the cost higher education has not yet priced.

That is the argument of this paper, and it is deliberately narrower than much of what is being written about artificial intelligence and higher education. We are not forecasting the end of the university. We are not claiming that AI will replace instruction. We are making a claim about a line in the operating budget that does not yet exist at most institutions, and about what happens to the value of the credential if it never does.

The prevailing debate has largely framed AI as a question of technology, of academic integrity, of pedagogy, or of workforce preparation. All matter. But before any institution can answer what students should learn *with* AI, it must answer a simpler question: how does an institution know what a student can do *without* it?

That question has always existed. AI has simply made it impossible to ignore. For decades, higher education operated on a practical assumption: work submitted by students was imperfect — sometimes original, sometimes assisted, occasionally dishonest — but generally close enough to individual capability that the distinction rarely threatened the value of the credential itself. Generative AI changes that assumption. For the first time, the appearance of competent work can be produced at almost no marginal cost. A student may submit work that appears sophisticated, organized, and professionally presented while revealing little about the capability of the individual who submitted it.

The problem this creates is not primarily moral. It is informational. And because credentials derive value from information, informational problems eventually become economic ones.

2. What a Degree Actually Sells

A credential is priced on the information it carries.[1] Its market value reflects what an employer can infer about the holder without conducting an independent assessment. That is the service being purchased, and it is why a degree commands a wage premium that a record of completed reading would not.

That inference depends on a mechanism: somewhere in the process, the institution observed what the student could do unaided and attested to it. When the mechanism works, the credential carries information. When it stops working, the credential stops carrying information, whether or not any individual student did anything wrong, because employers cannot distinguish between graduates for whom the mechanism still held and those for whom it did not.

[1] Michael Spence, "Job Market Signaling," *Quarterly Journal of Economics* 87, no. 3 (1973): 355–374; developed at length in Spence, *Market Signaling: Informational Transfer in Hiring and Related Screening Processes* (Harvard University Press, 1974). The informational value of credentials and the consequences of imperfect signaling underpin the framework used throughout this paper.

This is why treating AI primarily as a conduct problem has been a category error. Academic integrity asks who cheated. Verification asks what the institution can still attest to. By verification, we mean an institution's ability to credibly attest that a student can perform without assistance at the level the credential implies. An institution can maintain an exemplary academic-integrity record while degrading the informational value of its credentials. The first is a compliance question. The second is a business question. The labor market ultimately cares about the second.

The distinction is increasingly visible in the language of the major university reviews. Cornell's Committee on the Future of the American University observes that AI "has already proven adept at mimicking the form of knowledge without the supporting substance of inquiry."[2] That matters because employers rarely pay for the diploma. They pay for the capability the diploma is supposed to represent.

The detection route, meanwhile, is closing. MIT's faculty committee recommends against reliance on AI detectors, citing an escalating arms race, the risk of false positives, and harm to the instructional relationship; MIT's disciplinary guidance does not accept detector output alone as evidence.[3] Institutions that purchased detection software purchased time, not verification.

And the market that prices credentials may already be responding. New-graduate unemployment reached 5.6 percent in early 2026, up roughly 1.6 percentage points over three years, with employment declines concentrated in several entry-level occupations even as experienced workers held comparatively steady. The causal role of AI remains contested, and the literature is careful on the point: Stanford researchers found unemployment among the most AI-exposed workers rose 0.77 percentage points since 2022 against 0.85 points among the least exposed, and note that some of the shift preceded ChatGPT.[4] This paper takes no position on whether AI is already driving entry-level labor-market weakness. It does not need to. The important observation is simpler: employers will respond to changes in what credentials tell them long before accreditors issue guidance or universities complete curriculum reviews. If confidence in credential information declines, labor markets will react first. Universities will discover the change later.

The problem, in short, is not that students can use AI. The problem is that institutions must still be able to certify what students can do when they do not.

3. The Line That Matters Is Not the Department. It Is the Line of Sight.

The sector-wide numbers conceal the most useful fact available to a leadership team.

Nationally, 74 percent of faculty report that their students use AI to write essays, and 46 percent believe half or more do. But those averages span an enormous range. In English, history, the humanities, and

[2] *Report of the Committee on the Future of the American University*, Cornell University, September 22, 2026.

[3] *Report of MIT's Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training*, Massachusetts Institute of Technology, dated August 13, 2026, released September 2026.

[4] Neale Mahoney, Erika McEntarfer, and Karsen Wahal, "What Is Really Happening to Jobs? Separating AI Hype from Reality," Stanford Institute for Economic Policy Research policy brief, July 2026. Both the new-graduate figures and the exposure comparison are from this brief.

the social sciences, at least 88 percent of faculty report students using AI for essay writing, and 60 percent of English faculty report moderate or significant challenges managing it. In mathematics, 39 percent of faculty report no significant challenges at all, with engineering and the health sciences close behind.[5]

The obvious reading is that this is a humanities problem. It is not, and the distinction determines what an institution should do next.

Mathematics is not safer because it is mathematics. It is safer because its assessments are typically supervised. A proctored examination, a laboratory practical, a clinical rotation, a studio critique, and an oral defense all share one property: the institution watches the student perform. A take-home essay does not have that property. Neither does an unsupervised coding project, a remote quiz, or a paper written over a weekend. The exposure follows the assessment design, not the discipline.

That reframing matters more than the statistics, because it converts a sector-wide observation into an audit any provost can run. For every required course in every degree program, ask one question: **is the graded work produced where someone can see it?** The answer creates a map of institutional exposure, and it will not resemble the organizational chart. There are humanities sequences that are largely protected. There are technical programs that remain substantially exposed. There are professional programs that verify nearly every meaningful competency, and programs traditionally considered rigorous whose critical assessments now occur almost entirely out of sight.

The moment an institution asks the question this way, the problem becomes operational rather than ideological. It becomes possible to identify where verification remains intact, and where it has quietly disappeared.

4. Where Verification Still Works

One advantage of viewing the issue through assessment design rather than discipline is that parts of the solution are already visible. Institutions do not need to invent verification from scratch. They need to recognize where it has continued to function.

MIT's committee is explicit that learning goals involving laboratories, projects, and performance-centered activities may require relatively little adjustment in response to AI.[6] The national faculty data points the same direction: engineering, mathematics, and health-science faculty report the fewest difficulties, not because their students refrain from using AI, but because supervised practice, demonstration, critique, and in-person evaluation never depended heavily on artifacts produced out of sight.[7]

---

[5] Emily Angehr, Michael Bloem, Jessica Howell, and Alexandria Walton Radford, *College Faculty Perceptions of Generative Artificial Intelligence in Higher Education*, College Board Research, February 2026. Survey of more than 3,000 faculty at roughly 1,000 institutions, fielded June–August 2025.

[6] See note 3.

[7] See note 5.

Two implications follow. The first is institutional: programs built around supervised demonstration are not peripheral examples. They are working reference models for what verified assessment looks like in practice. In a conversation dominated by concerns about essays, these fields have been treated as exceptions. They may be closer to the solution than anything else on campus.

The second is pedagogical: for these programs, AI is often less a threat than an amplifier. MIT reports that AI can expand the ambition of project-based learning, allowing students to pursue work that would previously have been impractical within a semester. The challenge, in other words, is not AI itself. It is verifying who can perform. Where verification remains strong, AI can be incorporated without threatening the informational value of the credential. Where verification has weakened, AI exposes the weakness.

5. What Restoration Costs

Two new recurring costs arrive at the same time. Most institutional planning has recognized one of them.

The technology line

Providing AI access to students and faculty is rapidly becoming a condition of operating, and it is an operating expense rather than a capital one.

The cost is increasingly straightforward to estimate. The California State University system provides ChatGPT Edu to roughly 533,000 students, faculty, and staff for about $13 million a year, near $24 per user per year, and close to the floor of what is achievable at that scale.[8] At the other end of the range, MIT provides its community roughly $30 per user per month in credits through a model-agnostic platform, and its own committee reports that this "may not be sufficient," with some groups still paying for commercial subscriptions.[9] At an institution of 20,000 students, the CSU rate implies annual spending under $500,000; provisioning at MIT's level implies several million dollars a year. What separates them is negotiating position and how much actual usage the institution enables.

The key point is not the amount, but the shape of the expense. Technology spending has the features boards and finance committees often favor: it is visible, measurable, and contractable, and it can show results within a single semester through an announced enterprise agreement, provisioned accounts, and reported participation metrics. That is not a criticism; many institutions should make these investments. The problem is that AI access and verification are different challenges. An institution can become highly effective at teaching AI while remaining uncertain about what students can do without it.

---

[8] Kate Rix, "California State University renews controversial systemwide contract with OpenAI," *EdSource*, May 21, 2026; contract figures confirmed by a CSU spokesperson. The original 18-month agreement was reported by Angel Corzo, *CalMatters*, May 1, 2026, from the contract obtained under public records law; neither CSU nor OpenAI publishes the contract value. The tiered structure of the original agreement—a 40,000-user pilot followed by full-scale deployment—was reported by Julia Barajas, *LAist*, August 19, 2025.

[9] See note 3.

There is also an equity dimension that deserves more attention than it receives. MIT reports commercial subscriptions running well above the level of institutionally provided credits, and warns of large performance differences between students who can supplement institutional access with personal spending and students who cannot. A capability gap like that inside a single course is a form of educational inequality no financial aid formula currently contemplates.

Technology access, then, is undeniably a new cost. It is simply not the cost this paper is primarily concerned with.

The instructional line

Restoring verification where it has weakened means moving assessment back under observation. That is not a software purchase. It is oral examinations, supervised writing, portfolio defense, performance assessment, structured conversation and, ultimately, people.

The challenge becomes apparent whenever institutions move from principles to implementation. After five months studying its own campus, MIT's committee recommends reviewing learning goals across the curriculum, redesigning assessment, increasing instructional support, adapting learning spaces, creating departmental AI leadership, funding experimentation, expanding professional development, establishing new metrics, and considering limits on class size.[10] That is not a technology roadmap. It is an instructional roadmap, and it carries instructional costs. The report is particularly revealing because it comes from an institution with extraordinary resources and one of the deepest AI faculty benches in the world, which still notes that its existing teaching-support resources are already operating at capacity.

The shape of the cost is the problem. Most forms of credible verification scale with contact hours, which means they scale inversely with section size. A ten-minute oral defense cannot be evaluated at the scale of a machine-graded assignment. Portfolio assessments require reviewers; performance demonstrations require observers; conversation-based evaluation requires time. The more confidently an institution wishes to certify capability, the more labor-intensive that certification becomes.

Technology often benefits from economies of scale. Verification frequently does not. That is the economic tension at the center of this paper: the interventions most likely to restore confidence in credentials are the interventions most likely to raise instructional cost per student, at institutions where section size is one of the few remaining levers holding unit economics together.

There is a second return on this spending that rarely appears in formal analyses. MIT documents declining office-hours attendance, reduced discussion participation, and fewer in-person study interactions during the first years of widespread AI adoption, and 84 percent of faculty nationally report that AI makes students less likely to engage deeply with course material.[11] The contact hours that restore verification are the same contact hours that hold the learning community together. A defended

[10] See note 3.

[11] See note 5.

portfolio simultaneously measures competence and deepens learning. For a residential institution, the same dollar defends the core asset twice.

Why only one of them will be funded

The two costs are not equally fundable, and the asymmetry has little to do with their relative importance. AI access appears in budgets as a vendor contract; verification appears as people. One can be announced; the other must be staffed. One is visible within months; the other may take years before its effects are evident. One resides in technology planning; the other reaches into academic departments, faculty workload, physical space, governance, and instructional budgets.

The result is predictable. **An institution that funds the first and defers the second will be able to describe itself as having an AI strategy. It will not have restored what its degree certifies.** That is the central prediction of this paper.

Even the most heavily invested institutions show the distinction. The University of Florida provides a useful illustration: beginning years before the public release of ChatGPT, UF invested substantially in AI infrastructure, faculty hiring, and curriculum development — one of the most ambitious institutional AI initiatives in American higher education.[12] Yet AI adoption and verification remain different challenges. Expanding AI instruction answers one question: how should students learn to use AI? Verification answers another: how does the institution know what students can do without it? Both questions matter. Only one directly determines the informational value of the credential.

6. Who Can Pay

Institutions know their own finances better than we do. The purpose of this section is not to diagnose the financial condition of higher education but to establish a narrower point: which institutions have room for two new recurring costs, and which do not. The costs associated with verification are recurring because they persist as long as the institution wishes to certify capability, and that makes the sector's existing financial position relevant.

[12] Lisa Eadicicco, "This university built an AI curriculum before ChatGPT. Now it wants to help other schools do the same," CNN Business, September 2, 2026; University of Florida news release, July 21, 2020, announcing the NVIDIA partnership and the commitment to hire 100 AI faculty.

Where the sector sits entering FY2027

Credit. Heading into 2026, all three major credit rating agencies published pessimistic sector assessments: Moody's and S&P each assigned a negative outlook, and Fitch labeled the sector "deteriorating."[13] None described the sector as uniformly distressed; the gap between financially strong and financially constrained institutions continues to widen. Fitch's FY2025 medians make the same point: a median adjusted operating margin of –0.5 percent across its 56 rated private institutions, with the highest-rated institutions far above that level.[14]

The budget offices. In the *Chronicle of Higher Education's* survey of administrators who work closely on college finances, only 20 percent said their institution's budget was sufficient to achieve its mission.[15]

Liquidity. Shulman and Horn found 15 of 44 midsize tuition-dependent New England private colleges facing serious liquidity challenges within three years at current enrollment; subsequent national analysis by Shulman suggests similar pressures extend well beyond the region.[16]

Pricing. The preliminary first-year tuition discount rate for 2025–26 reached 57.1 percent, up from a final 54.5 percent in 2024–25, and roughly one dollar in six of institutional aid is funded by endowment or gifts.[17] Meanwhile, the average net price at private nonprofits (approximately $32,830) has been roughly flat after inflation for twenty years.[18] Institutions have spent two decades preserving enrollment while absorbing growing financial pressure.

None of this demonstrates crisis. It demonstrates constraint. And constraint changes which solutions are financially and politically realistic.

---

[13] Fitch Ratings, *U.S. Public Finance Higher Education Outlook 2026*, December 4, 2025 (sector outlook: deteriorating); Moody's Ratings 2026 higher education outlook, November 2025 (negative), as reported by *Higher Ed Dive*, November 21, 2025, and *Inside Higher Ed*, November 25, 2025; S&P Global Ratings, *U.S. Not-For-Profit Higher Education 2026 Outlook: Lower Expectations for Higher Education*, December 2, 2025 (sector view: negative, a change from S&P's bifurcated 2025 view).

[14] Fitch Ratings, *Fiscal 2025 Median Ratios for US Not-for-Profit Private Colleges & Universities*, June 23, 2026; figures as reported by *Higher Ed Dive*, June 24, 2026. The margin includes endowment income; the FY2023–FY2025 trend is –1.7, –2.0, and –0.5 percent.

[15] The Editors, "How Bad Is Higher Ed's Budget Crunch?," *The Chronicle of Higher Education*, May 12, 2026. The underlying survey is the Chronicle's own, fielded fall 2025 (n = 275), published as *The Budget Crisis Is Here*, Chronicle Intelligence, February 2026.

[16] Steven M. Shulman and Michael B. Horn, "The Financial Risk of Declining Enrollment to Midsize Colleges," *Education Next*, November 26, 2025. The national extension is Shulman's *National Survey of Private Mainstream Colleges and Universities* (February 2026), as reported by Michael B. Horn in the *Chronicle* forum cited above.

[17] NACUBO, 2025 Tuition Discounting Study, results released June 1, 2026; figures as reported by *Inside Higher Ed* and *Higher Ed Dive*. In the 2024–25 data, endowment earnings (11.5 percent) and philanthropy (5 percent) fund about one-sixth of institutional aid; roughly half of institutional aid has no dedicated funding source.

[18] MIT Figure CP-10; also cited in the Cornell report. The figure is average net tuition, fees, housing, and food for first-time full-time students at private nonprofit four-year institutions, essentially unchanged in constant dollars from 2006–07 ($32,890) to 2025–26 ($32,830).

The budgeting problem hidden inside the problem

Most institutions already possess mechanisms for understanding financial risk: multi-year projections, enrollment forecasts, scenario analyses, liquidity planning. The issue is not whether planning exists. The issue is what assumptions it contains.

Technology expenditures are comparatively easy to insert into projections: a CIO can estimate software costs, procurement can negotiate contracts, usage can be modeled. Verification expenditures are different. They emerge through additional instructional labor, redesigned assessment, reduced section size, expanded teaching support, and ongoing professional development; these costs are difficult to isolate because they are embedded within ordinary educational activity. That makes them easier to postpone. And postponement creates an illusion: an institution that funds AI access but delays assessment redesign can appear responsive while leaving the verification problem untouched. The budget includes activity. The credential does not necessarily gain certainty.

One question is therefore worth asking of any multi-year plan: does it contain a line for both of the costs described above? If it contains only the technology line, it is funding the half that does not restore the credential.

7. What Money Alone Would Not Fix

Fully fund a typical institution tomorrow and four constraints would remain.

The failure mode is unmeasured. Nothing on a board dashboard detects students completing successfully while learning less. Grades hold steady or rise while capability erodes. There is no line item, no ratio, and no covenant that moves. The damage becomes visible first in employer hiring behavior, years after the decisions that caused it.

The incentive structure punishes the work. Restoring verification asks faculty to rebuild assessment across every course, under time pressure. Teaching redesign is among the least career-rewarded activities in the academy: largely invisible in promotion and tenure, and at research institutions in direct tension with what is rewarded. Nor is resistance the main obstacle; 69 percent of faculty who have never used AI hold negative views of it, against 38 percent of those who have, and nearly a quarter have never used it at all.[19] Unfamiliarity and misaligned incentives explain most of what looks like obstruction. Any implementation plan that assumes faculty will absorb this as a professional obligation is assuming away the binding constraint.

The burden falls on the units with the least standing. Seventy percent of English faculty have a formal classroom AI policy against 34 percent in mathematics; negative sentiment runs from 66 percent in English to 25 percent in business and communications.[20] Any campus-wide committee is therefore a

[19] See note 5.

[20] See note 5.

body in which most members are not experiencing the emergency the minority is, and the departments carrying it are often the ones that have been shrinking and losing budgets.

And rigor carries competitive risk. An institution that moves assessment back under observation, while the institution twenty miles away does not, is taking an enrollment risk. At a discount rate near 57 percent, in a buyer's market, during a demographic decline, that risk is not hypothetical. Unilateral reform is individually irrational even where it is collectively necessary, which is why this is a problem for consortia, accreditors, and disciplinary societies rather than for any single institution acting alone.

8. Divergence, Not Convergence

The tendency in higher education is to assume that major shifts eventually diffuse across the sector: a practice originates at a handful of institutions, spreads through colleges and associations, appears in accreditation language, and becomes normal. That history encourages a comforting assumption—if the pedagogical answer becomes clear, the sector will converge around it.

The evidence suggests a different possibility. The sector may not converge. It may diverge. The reason is straightforward: the emerging consensus around AI does not primarily require better ideas. It requires resources, and institutions are entering this transition with radically different capacities. An enterprise AI license can be purchased in a procurement cycle. Verification cannot. It must be designed, staffed, managed, and sustained.

The distribution of that capacity is the troubling part. The institutions most able to fund verification restoration are often under the least pressure to do so: they hold stronger balance sheets, greater staffing flexibility, and — critically — already rely heavily on forms of supervised assessment, with smaller sections and stronger residential experiences. The institutions facing the greatest verification challenge are often tuition-dependent, with larger sections, heavier enrollment dependence, and less financial flexibility. They may need restoration more urgently while possessing fewer resources with which to pursue it. The pedagogical answer may be broadly shared; the financial ability to implement it is not.

If that asymmetry persists, AI introduces a new axis of stratification: not by prestige, but by verification. Prestige describes how institutions are perceived. Verification describes how confidently an institution can demonstrate capability. For decades the two moved together closely enough that the distinction rarely mattered. In the AI era it may matter a great deal, because verification directly affects the informational quality of the credential itself. Two institutions with similar missions and comparable programs could come to differ substantially in their ability to demonstrate what their graduates can do. Both issue degrees. The informational content of those degrees would no longer be equivalent. And in information markets, small differences in confidence produce meaningful differences in value: generative AI makes the *appearance* of capability abundant, which makes *evidence* of capability scarce – and scarcity creates value.

The most consequential divide in higher education over the next decade may therefore not be between institutions that adopt AI and institutions that resist it. Most institutions will adopt AI. The divide will

emerge between institutions that can restore verification and institutions that cannot. The irony is that both groups may describe themselves as following the same educational philosophy. Both may teach AI literacy. Both may endorse learning with and without AI. **The difference will not be what they believe. The difference will be what they can afford to implement.**

Divergence, we should be clear, is a risk rather than a destiny. Accreditors could put weight on demonstrated assessment integrity. Disciplinary societies could develop shared verification frameworks. Consortia could distribute the costs of experimentation, and assessment technology may improve. All of those deserve serious attention. But divergence becomes increasingly plausible when the solution is known, the solution is resource-intensive, resources are unevenly distributed, and implementation remains voluntary. Those conditions now exist.

9. The Employer's Side of the Ledger

The argument developed so far begins inside the institution. Its consequences appear outside it.

Employers do not experience AI as a curriculum problem. They experience it as an information problem. A hiring manager does not need to decide whether students should use generative AI in a writing course. A hiring manager needs to decide whether a candidate can write. Universities are concerned with how learning occurs; employers are concerned with what capabilities exist when learning is complete. For decades, the degree connected those two concerns efficiently: it reduced uncertainty, and that informational function is what gives credentials economic value. This observation predates AI by half a century. Spence's work on market signaling rests on the same principle: when direct observation is expensive, markets rely on credible signals as substitutes for complete information.[21] AI does not invalidate that framework. It intensifies it.

One misconception deserves to be addressed directly: credential erosion does not require employers to stop valuing degrees. In most cases they will not. The more likely response is supplementation. When confidence in a signal declines, markets rarely abandon it. Instead, they add signals. Technical interviews lengthen. Writing exercises, work simulations, and portfolio reviews become routine. Internships and certifications gain weight. The degree remains relevant; it simply becomes less sufficient. And the moment additional evidence becomes routinely necessary, part of the informational value of the credential has already migrated elsewhere.

This is what makes the erosion difficult for institutions to detect. Enrollment may hold. Graduation rates may hold. Accreditation reviews may remain positive. Nothing inside the university signals a problem, because the shift occurs in labor markets: employers quietly adjust, screening expands, and the institution notices years later, after the market has adapted. The consequences appear furthest from the decisions that cause them.

The strategic implication runs the other way too. If verification becomes scarce, institutions able to provide it acquire a differentiating asset, but credibility must be built. It emerges from assessment

[21] See note 1.

design, faculty practice, staffing, and institutional culture working together over time. An institution that can answer *how do you know?* — how do you know the student wrote that, how do you know the graduate can do this — possesses something increasingly scarce, and the more capable AI becomes, the more valuable that answer gets.

10. What Leading Institutions Have Actually Done

Over little more than a year, five major institutional inquiries appeared: MIT's ad hoc committee report, Cornell's Committee on the Future of the American University, the University of Chicago's working-group report and presidential statements, the AI at Duke Steering Committee's report to the provost, and a white paper on assessment from Stanford's Accelerator for Learning in collaboration with ETS, one of the world's largest educational testing organizations.[22] Together they represent hundreds of interviews, meetings, and institutional deliberations. Reading them together is instructive in two directions.

They converge on the pedagogical answer. The most succinct formulation comes from the University of Chicago's president, Paul Alivisatos, who argues that students should learn to think *with*, *without*, and *about* AI, sequenced so that machine assistance follows "a strong foundation of human-centered thinking and learning in community."[23] MIT arrives at a similar destination using different language: effective, responsible, and ethical AI use as distinct but interrelated goals, a principle of augmentation rather than automation, and explicit recommendations for deliberately AI-free experiences.[24] Cornell refuses the same binary in nearly the same words: the response should be "neither to restrict AI nor to embrace it uncritically, but to become far more intentional about what remains indispensably human in the learning process."[25] Stanford and ETS focus on the assessment half, arguing evaluation must move toward portfolios, performance tasks, and richer evidence of learning.[26] Duke, whose work emphasizes literacy, ethics, and institutional capability, broadly aligns.[27] Across institutions with very different cultures and governance, one pattern is emerging: students must learn to think without AI before they can effectively think with it.

They diverge on implementation. This matters because universities do not grant degrees on principles: they grant them on requirements. MIT recommends. Stanford advises. Duke builds capability. Cornell articulates a vision and explicitly declines a unified core curriculum. Chicago is different: it has altered requirements. A new Core writing course, *Inquiry, Conversation, Argument*, is

[22] The five reports are cited individually in the notes to this section and listed in full under Sources.

[23] Paul Alivisatos, "Teach Students to Think With, Without and About AI," *Wall Street Journal*, September 20, 2026.

[24] See note 3.

[25] See note 2.

[26] Nneka J. McGee, Candace Thille, Ikkyu Choi, Kadriye Ercikan, and Isabelle C. Hau, *Responsible Assessment in the AI Era: Key Insights from a Future-Focused Convening*, Stanford Accelerator for Learning, Stanford University, in partnership with ETS, 2026. Candace Thille, who led the Stanford effort, serves on the ETS board of trustees.

[27] AI at Duke Steering Committee, *Report to the Duke University Provost*, public summary, February 2026 (released March 2026).

the first change to Chicago's Core writing requirement in more than 25 years; it is being phased in now and reaches every first-year student by 2027–28. The university does not describe it as an AI ban; Alivisatos has publicly described it in terms of real writing that begins with a blank page and sustained engagement with instructors and peers. In addition, its Social Sciences Collegiate Division has adopted restrictions on devices and AI in its core sequence, and its Law School's first-year curriculum deliberately establishes writing and reasoning without AI before layering AI-assisted work on top.[28] The significance is not the specifics of Chicago's policies. The significance is that Chicago changed requirements. A university can publish a report and change nothing. A requirement changes behavior.

And none of them has solved the cost. No major report claims verification is inexpensive. None proposes a mechanism that restores confidence in assessment without additional instructional investment. None argues AI has made supervised assessment less valuable, and if anything, they move in the opposite direction. None of these institutions is resource-constrained, which means the implementation gap at the top of the sector is not primarily a money problem, and the distribution problem below the top is worse than the financial analysis alone suggests. The sector is converging on the educational answer. The remaining challenge is economic.

11. The Objection

The strongest objection to this paper does not come from the institutions resisting change. It comes from the institutions leading it, and no one has articulated it more clearly than Alivisatos, who opens his own argument by setting aside the familiar worries (pervasive cheating, students abandoning college, credentialism) as concerns that, while real, miss the larger story about **how to educate students for a world abundant with AI**.[29]

It is a serious argument, and it deserves an answer in three parts.

First: he is right about the pedagogy, and we say so without hedging. The with and without framework, and the sequencing that follows from it, is the best formulation anyone has produced. Students cannot effectively delegate judgment they do not possess, cannot critically evaluate outputs they do not understand, and cannot responsibly use tools whose limitations they cannot recognize. It is not a compromise between the prohibitionists and the enthusiasts. It is a better answer than either.

Second: we are answering a different question. Alivisatos asks *what should students learn?* This paper asks *what does it cost to know that they learned it, and who can pay?* One question does not answer the other. An institution may possess a brilliant educational philosophy and still lack a mechanism for verifying outcomes; higher education has encountered that gap before. His question is the more

[28] *Rethinking Legal Education in the AI Era*, University of Chicago Law School, July 9, 2026. The Social Sciences Collegiate Division policy, effective fall 2026, was issued by divisional memo and is documented in Shea Vance, "U. of Chicago's Social Sciences Division Bans Tech, AI in Core Courses," *The Chronicle of Higher Education*, August 24, 2026. The *Inquiry, Conversation, Argument* requirement was announced in University of Chicago News, December 11, 2025.

[29] See note 23.

important one intellectually. Ours is the one that determines which institutions are in a position to act on his.

Third: the concession. If every institution could do what Chicago has done, this paper's argument would be wrong, and we would be glad. Suppose every faculty senate endorsed the framework tomorrow. What happens next is implementation, and implementation is not distributed evenly. Some institutions can absorb increased instructional costs and others cannot. Agreement does not eliminate divergence. In some cases, it makes divergence more visible. The claim is not that the problem is unsolvable. It is that **the answer is known, cheap to state, expensive to staff, and distributed in inverse proportion to need**.

12. Strategic Implications

Artificial intelligence did not create higher education's verification problem. It exposed it. And the exposure changes how institutions should think about strategy itself, because the strategic question is not which AI tools to adopt. The strategic question is what the institution sells. Universities hold no monopoly on information — a student can access more of it in an hour than earlier generations could in months. The institution's distinctive value lies in judgment and verification: *we observed this person learning; we evaluated that learning; we can attest to what this person can do*. That capacity is the foundation on which the credential's value rests, and it is precisely the capacity AI forces institutions to examine.

Assessment has traditionally been treated as a pedagogical concern: faculty discuss it, teaching centers support it, accreditors review it, and boards rarely think about it. That may no longer be sustainable. Assessment drives verification: verification drives informational quality; informational quality drives signaling value; signaling value shapes employer confidence; and employer confidence sets the economic position of the credential. The chain is longer than most strategy discussions acknowledge, but once visible it is difficult to ignore. Institutions that restore verification are not merely redesigning courses. They are protecting the informational value of their product.

Three asymmetries define the landscape. **Timing:** AI evolves at the speed of capability release; institutions evolve at the speed of governance. New capabilities arrive on a cadence measured in months; curricular revision takes years. Nonetheless, Chicago's reforms show the constraint is not a law of nature, compressing what its president notes would once have been a decade-long debate into roughly two years.[30] The argument does not depend on that pace continuing. The verification problem is already created by capabilities that have shipped. **Cost:** two recurring expenditures arrive together, one a vendor contract and one people, competing for the same resources, and only one currently receives sustained attention. **Visibility:** technology investments produce immediate evidence; an improved oral examination process rarely generates headlines, and its absence produces nothing a board can see either, until employers notice, which is too late to be useful.

[30] See note 23.

Together, the asymmetries create the paper's most important caution: **the risk of mistaking activity for progress**. AI centers can be launched, policies written, and licenses acquired. All are worthwhile, though none of them automatically restores verification. Institutions measure what they do rather than what they certify. The relevant question is not *what AI initiatives have been announced?* It is *what evidence now exists that did not exist before?* Evidence of capability. Evidence of performance. Evidence that the credential continues to represent something real. Many of the decisions shaping future verification capacity (class sizes, faculty hiring, teaching-support staffing, learning-space design, assessment policy) are being made now, by leaders who are not describing them in those terms.

13. What to Watch

If this argument is correct, the most important developments over the next several years will be institutional rather than technological, and they should be observable. One caveat first: institutions have had solvency instrumentation for decades and closed anyway. Metrics are necessary and plainly not sufficient. What follows is not a dashboard. These are signals worth watching.

Does instructional cost per student move? If verification restoration is genuinely labor-intensive, it must appear somewhere. Oral examinations, portfolio reviews, and supervised assessment all require people. If institutional AI spending rises while instructional effort stays flat, the technology line is being funded and the verification line is not. One expenditure provides access; the other provides evidence, and evidence is the scarce asset.

Do requirements change, or only recommendations? MIT recommends; Cornell articulates principles; Stanford and ETS advise on assessment; Chicago has altered degree requirements. Recommendations reveal institutional understanding. Requirements reveal institutional commitment. The strongest signal that the sector has entered a new phase will not be another report. It will be evidence that institutions are changing what students must do in order to graduate, and how quickly.

Does employer behavior shift? The leading indicator of credential strain is hiring practice, not enrollment. Watch for expanded technical assessments, writing evaluations, work simulations, portfolio requirements, and structured screening. None individually proves erosion; collectively they indicate rising demand for independently verified capability. Employers experience uncertainty before institutions do.

Does the line of sight move? The simplest measure any institution can run on itself: for every required course, what percentage of graded work is produced under observation? Every institution, department, and accreditation team can calculate it. If that share is not rising, then whatever else has changed (policies, centers, task forces, enterprise licenses) the fundamental exposure remains unchanged.

Does the student confidence gap close? In MIT's own data, roughly 70 percent of undergraduates believe AI skills will matter in their careers, while about 25 percent believe their institution is preparing

them to use AI in professional settings.[31] The percentages matter less than the question: students increasingly understand that AI matters, but do they believe their education is keeping up? A widening gap between perceived importance and perceived preparation is difficult to ignore, and any institution can measure its own in a single survey cycle.

These indicators do not measure understanding. They measure action. The next several years will reveal whether action catches up with understanding.

Conclusion

Three things are now true at once. The pedagogical answer is increasingly known, and the leading institutions have converged on versions of it. The cost of acting on that answer is concentrated in the one budget line — instructional labor — that the sector is least able to grow. And the failure mode is invisible: nothing on a dashboard moves while a credential quietly stops carrying the information employers pay for.

The task, then, is not discovery. It is restoration, and it is expensive in exactly the way institutions are worst equipped to fund: recurring, instructional, invisible in the near term, and needed most where it can be afforded least.

In a world where the appearance of competent work is free, the scarce thing is a credible claim that something is true: this graduate can actually do this. That is the business higher education has always been in. What this moment is revealing is how much of it had quietly come to depend on nobody checking.

Semper Verum Partners works with institutional leaders on strategy where the business model and the academic model must be reasoned about together. We can be reached at jj@semperverum.com or vw@semperverum.com.

[31] See note 3.

A Note on Method

Every figure in this paper traces to a named, published source, cited in the notes and listed below. Where a claim is contested in the literature, most notably AI's causal share of entry-level employment decline, we have said so rather than resolving it in our favor. Where a source could be verified only through secondary reporting, the notes say so.

In the spirit of the argument above: this paper was researched and assembled with the assistance of AI. The analysis, judgments, and conclusions are ours. A paper arguing that verification is the scarce good should say how it was made.

Sources

Primary institutional reports. Report of MIT's Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training (August 13, 2026) · Report of the Committee on the Future of the American University, Cornell University (September 22, 2026) · Nneka J. McGee, Candace Thille, Ikkyu Choi, Kadriye Ercikan, and Isabelle C. Hau, *Responsible Assessment in the AI Era: Key Insights from a Future-Focused Convening*, Stanford Accelerator for Learning, in partnership with ETS (2026) · AI at Duke Steering Committee, *Report to the Duke University Provost*, public summary (February 2026) · *Report from the Artificial Intelligence and Education Working Group*, University of Chicago Office of the Provost (July 31, 2025) · *Rethinking Legal Education in the AI Era*, University of Chicago Law School (July 9, 2026) · Paul Alivisatos, "Teach Students to Think With, Without and About AI," *Wall Street Journal* (September 20, 2026)

Quantitative sources. Emily Angehr, Michael Bloem, Jessica Howell, and Alexandria Walton Radford, *College Faculty Perceptions of Generative Artificial Intelligence in Higher Education*, College Board Research (February 2026) · Neale Mahoney, Erika McEntarfer, and Karsen Wahal, "What Is Really Happening to Jobs? Separating AI Hype from Reality," Stanford Institute for Economic Policy Research (July 2026) · Fitch Ratings, *U.S. Public Finance Higher Education Outlook 2026* (December 4, 2025) and *Fiscal 2025 Median Ratios for US Not-for-Profit Private Colleges & Universities* (June 23, 2026) · Moody's Ratings, 2026 outlook for US higher education (November 2025) · S&P Global Ratings, *U.S. Not-For-Profit Higher Education 2026 Outlook: Lower Expectations for Higher Education* (December 2, 2025) · NACUBO, 2025 Tuition Discounting Study (June 2026) · College Board, *Trends in College Pricing and Student Aid 2025* · The Editors, "How Bad Is Higher Ed's Budget Crunch?," *The Chronicle of Higher Education* (May 12, 2026) · Chronicle Intelligence, *The Budget Crisis Is Here* (February 2026) · Steven M. Shulman and Michael B. Horn, "The Financial Risk of Declining Enrollment to Midsize Colleges," *Education Next* (November 26, 2025) · Steven M. Shulman, *National Survey of Private Mainstream Colleges and Universities* (February 2026) · Michael Spence, "Job Market Signaling," *Quarterly Journal of Economics* 87, no. 3 (1973)

Reporting. Lisa Eadicicco, CNN Business (September 2, 2026), on the University of Florida · University of Florida news release (July 21, 2020), on the NVIDIA partnership · Kate Rix, *EdSource* (May 21, 2026), Angel Corzo, *CalMatters* (May 1, 2026), and Julia Barajas, *LAist* (August 19, 2025), on

the California State University–OpenAI agreements · Shea Vance, *The Chronicle of Higher Education* (August 24, 2026), on the University of Chicago Social Sciences Collegiate Division policy · University of Chicago News (December 11, 2025), on the *Inquiry, Conversation, Argument* requirement · *Higher Ed Dive* (June 24, 2026), on Fitch's fiscal 2025 median ratios · *Inside Higher Ed* and *Higher Ed Dive* (June 2026), on the NACUBO Tuition Discounting Study · *Higher Ed Dive* (November 21, 2025) and *Inside Higher Ed* (November 25, 2025), on Moody's 2026 higher education outlook