The Degree That Cannot Be Verified: Institutional Credentials in the Age of Ghost Resumes

Weekly Analysis — https://ainews.social

A recruiter at a mid-sized firm now opens a hiring inbox to a peculiar archaeology. The cover letters are crisp; the resumes are tonally consistent; the LinkedIn pages have a faint sheen of polish that suggests competent self-presentation. Some of these candidates exist. Some of them, in the sense that matters to the labor market, do not — they are composites generated in seconds, applied at scale, sometimes harvested by bots that never intended a human to read the response. The recruiter cannot tell which is which from the documents alone. Neither, increasingly, can the university whose seal sits at the top of the transcript. This is the situation the phrase "ghost resume" names, and it is the situation the American university has been slow to recognize as an existential one.

For the better part of a century, the bachelor's degree has functioned less as a description of what a graduate knows than as a statistical promise about who they are: someone who could be admitted, who persisted, who completed the rituals on schedule. The transcript was a signal — a portable token of trustworthiness underwritten by the institution's reputation. That promise is now in trouble for two reasons at once. The first is that the artifacts on which the signal travels — the resume, the cover letter, the writing sample, the take-home essay, even, in some programs, the capstone — are now trivially generable. Stanford's [20] documents the speed at which generative systems have crossed thresholds in writing, coding, and reasoning that two years ago demarcated trained from untrained labor. The second is that the institution itself, asked whether its graduates can in fact do what the seal implies, has no good answer beyond the seal.

[20] The 2026 AI Index Report

This essay is about what happens to a higher education sector whose primary product is a credential that the labor market is starting to mistrust, and whose response to that mistrust has so far been to invest in surveillance rather than verification. The argument, in short, is that the ghost resume is not a hiring problem to be solved by HR departments armed with better detectors; it is a higher-education problem that forces the sector to choose between issuing proofs that

mean something and continuing to issue tokens that increasingly do not. The good news is that some of the work — capstones, apprenticeships, portfolios, oral defenses — is already familiar to faculty. The bad news is that almost none of the present discourse in U.S. higher education is organized around scaling that work. Most of it is organized around catching cheaters.

What the Diploma Was For

The signaling theory of education, much hated by educators and much loved by economists, holds that the degree's primary economic function is not the transfer of skill but the cheap broadcast of a candidate's underlying traits to employers. On this account, the institution is a filter, the diploma is the receipt, and the labor market trusts the receipt because it trusts the filter. The theory has always been partial — universities also teach things — but it described the transaction well enough that decades of policy were built on top of it.

What generative AI breaks is not the teaching part. It breaks the receipt. When [20] and survey work like the [17] show that the overwhelming majority of undergraduates are now using generative tools routinely in coursework — not as occasional shortcuts but as default infrastructure — the question of what an A in a writing-intensive course actually certifies becomes genuinely unclear. The HEPI survey, the most-cited longitudinal student instrument in the English-speaking world, finds usage approaching saturation across disciplines, with the largest gains among precisely the assignments transcripts were designed to evaluate.

[20] The 2026 AI Index Report

[17] Student Generative Artificial Intelligence Survey 2026

This is a problem that the labor market is not going to wait for higher education to solve in committee. A faculty member surveyed by Forbes' coverage of the [2] report can tell themselves a story about cognitive offloading and the erosion of foundational skills, and that story may well be true; but the story the recruiter tells herself is simpler and more damaging. It is that the transcript, the writing sample, and the senior thesis no longer reliably distinguish a graduate who can think from one who can prompt. If she is right, the institution's signal is already devalued, regardless of what the institution still believes it is selling.

[2] 90% Of Faculty Say AI Is Weakening Student Learning: How Higher Ed Can Reverse It

The [26] literature has been arguing for two years now that the artifact-based theory of assessment — the idea that we evaluate the document the student hands in — is finished, and that the only durable replacement evaluates the process that produced it. The institutional response to this argument has been, at best, uneven, and the

[26] Writing with machines? Reconceptualizing student work in the age of AI

unevenness is the substrate on which the ghost resume crisis grows.

Detection Theater

The first instinct, when a signal is corrupted, is to police the corruption. American universities, under pressure from faculty and trustees alike, have spent the past three years pouring money into AI-detection software that, by every credible audit, does not work. CalMatters' investigation, [8], documents a procurement landscape in which campuses have signed seven-figure contracts with vendors whose own published false-positive rates would, if applied at scale, generate hundreds of wrongful-conduct findings per term. The detectors, the reporting shows, systematically misclassify the writing of multilingual students and students with disabilities — a familiar pattern in algorithmic systems trained on a thin majoritarian baseline.

[8] Colleges pay millions for AI detectors that are flawed

The pedagogical and legal consequences are predictable. Newsday's account of the [4] case is one of a growing genre: student writes paper, instructor runs paper through detector, detector flags paper, student is sanctioned, student demonstrates that the detector is wrong, student sues. The institution, in these cases, finds itself defending a procurement decision rather than an educational principle, and the procurement decision is hard to defend because the underlying tool is hard to defend.

[4] Adelphi accused a student of using AI to plagiarize. He sued.

The deeper problem is not that the detectors are bad — although they are. It is that detection-centered governance encodes a theory of academic integrity in which the institution's job is to certify the absence of generative AI in student work, rather than to certify the presence of student understanding. That theory was barely coherent before tools became invisible inside the operating system; with AI assistance now embedded into Microsoft and Google's productivity suites and into the assistive-technology layer that universities are obligated to provide under federal law (a fact made vivid by Microsoft's own training module on [15]), the absence theory is incoherent on its face. There is no AI-free baseline to detect deviations from.

[15] Personnaliser l'apprentissage pour les étudiants handicapés à l'aide de l'IA

The detection-industrial complex is, in this light, a way for institutions to look busy while postponing the harder conversation. It substitutes surveillance for assessment redesign. It also drains money and trust at exactly the moment both are needed for the redesign work. The Australian higher-education commentariat has been more forthright about this than the American: a recent piece on [3] argues bluntly that framing AI as a cheating problem rather than an assessment-design problem has loaded the entire cost of institutional

[3] Academic Staff Are Paying The Price For The Misframed GenAI Assessment Debate

indecision onto teaching staff, who are now expected to police what their employers will not redesign.

What Students Are Actually Doing

If detection is a bad theory of where the threat lies, it is partly because the threat is misdescribed. The picture of the student as a strategic cheater — calculating whether to outsource an essay — is a comforting one for administrators because it locates the moral failure at the individual level and exempts the institution from rethinking the assignment. The empirical picture is messier and more interesting.

Inside Higher Ed's reporting on [23] describes a use pattern that does not map cleanly onto the cheating frame at all: students are using generative tools as study partners, as accessibility aids, as therapists, as drafting scaffolds, and — increasingly — as the only available office hours when faculty are stretched thin. The same story appears in the [17], where the modal student reports a relationship with these tools closer to dependence than transgression. Times Higher Education has been reporting the same finding from a different angle: when surveyed, [18] — they want to be taught how to use these tools well, and they perceive the absence of such teaching as institutional negligence.

[23] The Myriad Complex Ways Young People Use AI

[17] Student Generative Artificial Intelligence Survey 2026

[18] Students are asking for AI guidance, not just policy

This is a real generational fact, and it cuts in two directions. On one hand, it complicates the moral simplicity of the cheating frame: a student using a chatbot to clarify a confusing reading is not obviously doing the same thing as a student outsourcing a final paper, and treating the two identically is a failure of pedagogical seriousness. On the other hand, it makes the credential problem worse, not better, because if students are now thinking-through-AI as a default cognitive practice, the question of what the credential certifies — the student's mind, or the student's mind in collaboration with a tool whose capability changes monthly — has no stable answer. The Nature paper on [5], which finds that an AI tutor produced larger learning gains than an active-learning lecture in a controlled physics experiment, is an interesting result and a deeply unsettling one: it suggests that in some domains the most effective teacher in the room is no longer the human being whose name is on the syllabus and whose grade is on the transcript.

[5] AI tutoring outperforms in-class active learning: an RCT

A three-level meta-analysis published this year, [11], finds the picture more equivocal — generative AI yields measurable learning gains under some pedagogical conditions and measurable losses under others, with the determining variable being the structure of the task

[11] Exploring the effect of GenAI on learning outcomes in higher education: A three-level meta-analysis

rather than the tool. There is also serious work suggesting that heavy reliance erodes precisely the skills the credential claims to certify. [22] reports decrements on standard instruments among students who use AI as a substitute rather than a scaffold, and a working paper from the University of Chicago's data science institute, [24], argues that even the rhythm of AI access — whether students use it before, during, or after engaging a problem on their own — substantially shapes whether the technology trains or replaces the underlying cognition.

[22] The Impact of AI on Students' Reading, Critical Thinking, and Problem Solving

[24] The Time Constraints of AI Access Could Change How We Think

The takeaway here is not that AI is bad for learning, or good for it. It is that the credential — the diploma, the transcript, the GPA — was designed to summarize a set of cognitive achievements that AI is now reorganizing in ways the credential cannot represent. A 3.8 means something different in 2019 and in 2026, and the institution issuing it has not updated the meaning.

The Governance-Pedagogy Imbalance

Walk into a faculty senate on almost any campus and the AI agenda item is, with depressing reliability, a governance item: a proposed academic-integrity policy, a syllabus statement template, a list of approved and unapproved tools, a procurement update on the detection vendor. Walk into a curriculum committee in the same week and AI is, with comparable reliability, absent from the agenda entirely. The governance discourse has metastasized; the pedagogical one has barely begun.

This imbalance is visible in the labor market for academic positions itself. The University of Florida's [9] job posting — one of dozens like it across R1 institutions — describes a role focused on integrating AI into existing course architectures rather than rethinking what courses are for in the first place. The work is being done, in other words, by a contingent professional class of instructional designers, who have neither the standing nor the mandate to redesign assessment regimes, while the tenured faculty who do have that standing are largely insulated from the question.

[9] Details - Instructional Designer Campaign - University of Florida

The MDPI piece [13] names this imbalance directly: the policy infrastructure for AI in higher education is racing ahead of the curricular infrastructure, and the result is a sector that has many rules about what students may not do with AI and very few coherent ideas about what they should do with it. A companion paper, [7], and the practitioner-focused [6] report make the constructive case: the only assessment formats that survive the generative wave are those that watch the student think rather than collect what the stu-

[13] From Cognitive Necessity to Cognitive Choice: Higher Education Assessment and Learning in the Age of Generative AI

[7] Beyond Detection: Redesigning Authentic Assessment in an AI Era

[6] Authentic Assessment in the Age of AI

dent produces. Oral defenses, in-process critiques, supervised studio work, viva-style examinations, capstones with public presentation, apprenticeship-style portfolios that document iteration — all of these formats already exist in the academic toolkit, and all of them are expensive in the one resource the modern American university has stripped to the bone, which is faculty time.

This is where the misframing of the GenAI debate as primarily a governance problem becomes a category error with budgetary consequences. The institution that responds to the credential crisis by tightening its conduct code and buying more detectors is solving for the wrong variable. The institution that responds by reinvesting in low-faculty-ratio assessment formats is solving for the right one, but the bill is large and the political coalition for paying it does not yet exist on most campuses. UNESCO's [16] makes this point about the institutional level: high-stakes deployments of AI in education require independent validation mechanisms, not just policy statements, and most institutions have not built those mechanisms because they cost real money.

[16] AI competency framework for teachers

A useful counter-example, often cited in the discourse, is the Arizona State University experiment described in [14], which redesigns an introductory writing course around in-class drafting, structured peer review, and oral component grades — formats in which AI assistance is either harmless or transparent. The course works, and the writeup is candid about why most institutions cannot replicate it without staffing decisions they do not want to make.

[14] Outsmarting AI in the classroom

The Credential as Cryptographic Object

If the diploma can no longer credibly signal what its holder can do, the technical question is whether some other artifact can. Here the discourse turns, often too quickly, to blockchain-based micro-credentials and verifiable digital diplomas — cryptographically signed records, issued by the institution and held by the learner, that an employer can verify against a public ledger without phoning the registrar. The technology is real and the basic idea is sound. A degree that a third party can verify in seconds, that cannot be forged without breaking the underlying cryptography, and that the institution can revoke if it later determines the underlying work was fraudulent, is straightforwardly an improvement on the PDF transcript.

But the technology does not by itself solve the signal problem. A cryptographically signed credential certifying that the holder completed a writing-intensive seminar in spring 2026 is no more meaning-

ful than the PDF version if the writing was generated by a system the institution declined to ask about. The verification stack secures the chain of custody from the institution to the employer; it does not secure the chain of custody from the student to the institution. Without a redesigned assessment regime underneath, the cryptographic layer just makes ghost credentials more efficiently transmissible.

There is also a coordination problem that the discourse tends to underplay. A verifiable credential standard is only as useful as the network of institutions and employers that recognize it; a single university issuing blockchain diplomas into a market that still asks for the PDF transcript is solving a problem nobody in the transaction perceives. The Spanish-language strategic document [10] is one of the more lucid statements of the cross-institutional infrastructure this would require — interoperable issuance, shared revocation registries, common skill ontologies — and the gap between that document and the actual procurement state of any given American campus is enormous.

[10] ENIAG2025_PRINCIPIOS

There is a still deeper concern, which the literature on algorithmic systems has been articulating for years and which Ruha Benjamin's [16] frames most sharply: when verification infrastructure is built on top of a corrupted social baseline, it tends to launder that baseline rather than correct it. A skills-credentialing ecosystem that issues fine-grained micro-credentials in machine-readable form is also a surveillance ecosystem. It produces a permanent record of what a learner has and has not been certified to do, accessible to any employer with the right key. The Eubanks-Benjamin tradition has been documenting for a decade what happens when the labor market gains finer-grained data about the working class; the answer has not generally been more equitable hiring. The fact that ghost resumes are a real problem does not mean the solution is a system in which every learner walks around with a verifiable transcript-on-a-chain that follows them for life.

[16] Race After Technology

The Legitimacy Crisis Underneath

The ghost resume phenomenon arrives in a sector already under unusual strain. American higher education is in the middle of a multi-year crisis of public legitimacy, driven by some combination of cost, grade inflation, equity scandals, and a partisan campaign against the credentialing class. The credential problem is not the first thing eroding the diploma's signal value; it is the most recent and most technologically vivid in a sequence that includes the bipartisan critique of grade inflation, the post-affirmative-action litigation over admissions, and the long debate over whether the bachelor's degree is now

functioning primarily as a class marker rather than a competence one.

This matters for how the sector should think about its response. An institution facing only the AI problem could plausibly redesign assessment, ship verifiable credentials, and ride out the transition. An institution facing AI plus a legitimacy crisis cannot, because the redesign will be read — by trustees, legislatures, and editorial pages — through the lens of the legitimacy crisis. Every restructured assessment regime will be evaluated for whether it is "rigorous." Every micro-credential will be evaluated for whether it dilutes the bachelor's degree. Every move toward process-based assessment will be evaluated for whether it advantages students whose home environments support unstructured, iterative work — a longstanding equity concern that the authentic-assessment literature, including [7], takes seriously and the public discourse mostly does not.

[7] Beyond Detection: Redesigning Authentic Assessment in an AI Era

The piece [21] makes a pointed version of this argument: institutions are at risk of using AI to optimize the parts of higher education that are least defensible — the throughput, the standardized assignment, the LMS-mediated transaction — and to neglect the parts that are most defensible, namely the slow human relationships in which judgment is formed. If that warning is right, then the credentialing crisis is also a values test. An institution that uses AI to grade more papers faster is producing a credential whose value continues to fall. An institution that uses the credentialing crisis as the occasion to rebuild around the practices that AI cannot substitute for — close mentorship, oral defense, supervised practice, public-facing capstones — is producing a credential whose value might stabilize.

[21] The AI-Native University Must Guard Against Getting Better at the Wrong Things

The CORE working paper [25] makes a useful structural observation here: the institutions that will be hit hardest by the credential collapse are not the elite ones, whose signal is partly disconnected from what graduates can actually do, but the regional comprehensive and open-access institutions, whose graduates compete in labor markets where the credential is read literally. The legitimacy crisis is therefore not evenly distributed, and the sector's response so far — driven by the reputational anxieties of the most-protected institutions — has been correspondingly poorly calibrated.

[25] The Unintended Consequences of Artificial Intelligence and Education

What Survives

There is a thinner, sturdier version of the university hidden inside the present crisis, and most faculty already know what it looks like because they have been doing parts of it for years. It is a university in which assessment is dominated by formats that watch the student

think: studio critique, lab notebook, oral defense, juried portfolio, supervised clinical or field practice, apprenticeship-style work with documented iteration. It is a university in which the credential is not a single terminal artifact but an accumulating record of demonstrated competencies, each tied to a specific situated performance the institution can vouch for because a named human being watched it happen. It is a university whose use of AI in teaching — and there will be heavy use, because the [12] study and others show that faculty AI adoption is moving fast where institutional support exists — is structured around scaffolding student cognition rather than replacing it.

[12] Faculty Adoption of AI-Assisted Teaching Tools in Chinese Higher Education: An Integrated UTAUT2–TAM Framework at Shanghai University

What is missing from the present discourse, and what every honest reading of the evidence above points toward, is a partnership framing. The dominant U.S. institutional posture toward students on AI is adversarial: students are presumed to be cheating, faculty are deputized to detect it, vendors are paid to automate the detection, and conduct codes are sharpened to punish it. The dominant student posture, as the HEPI and Times Higher Education surveys show, is not adversarial but hungry — students want guidance, want sanctioned use, want to be taught the tool they are already using. There is a deal available here that almost no institution has made: explicit, generous permission to use AI for specified tasks, paired with assessment formats that require the student to demonstrate, in real time and under supervision, that they understand what the AI did. That deal would require faculty time, institutional courage, and a willingness to retire the artifact-as-evidence model that has organized higher-education assessment for a century. It is the deal the credential needs in order to mean something again.

The CORE compendium [19] collects faculty case studies from institutions that have struck versions of this deal, and the pattern is consistent: where AI use is named, scaffolded, and assessed in process, the credential survives the transition; where it is forbidden, hidden, and policed by detector, the credential degrades regardless of the conduct outcomes. The literature on critical AI literacy — the kind of work captured in resources like [1] — adds the layer that makes such a deal pedagogically defensible: students learning to interrogate model outputs, to identify failure modes, to recognize when the tool is confidently wrong, are doing intellectual work the credential can credibly certify, and are doing it on terrain where the AI cannot do the work for them.

[19] Teaching and Generative AI

[1] 24 Critical AI Literacy Questions Every Teacher Should Ask Students

The Commitment the Sector Has Not Yet Made

The argument of this essay has been that the ghost resume is a clarifying event for higher education, not because it introduces a new problem but because it makes a long-running one impossible to ignore. The diploma was always a thinner signal than the institution pretended. AI has not corrupted a previously robust system; it has made visible the system's existing thinness, and it has done so on a timeline shorter than the sector's ordinary capacity to respond.

The path that the discourse keeps gesturing toward — better detectors, sharper conduct codes, more aggressive procurement of integrity software — is a path that solves the problem on paper while making it worse in practice. It transfers the cost of institutional indecision onto faculty and onto wrongly accused students, and it postpones the redesign work while spending the budget that work would require. The CalMatters investigation into the [8] procurement landscape and the Adelphi litigation reported in [4] are not isolated incidents. They are the predictable output of the detection-centered theory.

[8] Colleges pay millions for AI detectors that are flawed
[4] Adelphi accused a student of using AI to plagiarize. He sued.

The path that the credential's survival actually requires is harder and more familiar than the technological discourse admits. It involves reinvesting in low-ratio assessment, naming and scaffolding student AI use rather than forbidding it, building cross-institutional verification infrastructure that the labor market can read, and accepting that the parts of the bachelor's degree that AI can fake are the parts that were never doing the certifying work to begin with. Stanford's [20] shows the capability curve continuing to steepen; the [17] shows usage already past the point at which artifact-based assessment is recoverable. The question is no longer whether the institution's primary product will be tested; the test is in progress, and the institution's grade so far is not what the seal would suggest.

[20] The 2026 AI Index Report
[17] Student Generative Artificial Intelligence Survey 2026

What the sector owes its graduates, and the labor market that hires them, is a credential that means something — not because it is wrapped in cryptography, although that helps, and not because it was produced under surveillance, which mostly does not, but because someone the institution pays watched the holder think, and is willing to put the institution's name on what they saw. That has always been what the diploma was supposed to certify. The ghost resume has merely reminded the sector that the certification is overdue.

References

1. 24 Critical AI Literacy Questions Every Teacher Should Ask Stu-

dents

2. 90% Of Faculty Say AI Is Weakening Student Learning: How Higher Ed Can Reverse It
3. Academic Staff Are Paying The Price For The Misframed GenAI Assessment Debate
4. Adelphi accused a student of using AI to plagiarize. He sued.
5. AI tutoring outperforms in-class active learning: an RCT
6. Authentic Assessment in the Age of AI
7. Beyond Detection: Redesigning Authentic Assessment in an AI Era
8. Colleges pay millions for AI detectors that are flawed
9. Details - Instructional Designer Campaign - University of Florida
10. ENIAG2025_PRINCIPIOS
11. Exploring the effect of GenAI on learning outcomes in higher education: A three-level meta-analysis
12. Faculty Adoption of AI-Assisted Teaching Tools in Chinese Higher Education: An Integrated UTAUT2–TAM Framework at Shanghai University
13. From Cognitive Necessity to Cognitive Choice: Higher Education Assessment and Learning in the Age of Generative AI
14. Outsmarting AI in the classroom
15. Personnaliser l'apprentissage pour les étudiants handicapés à l'aide de l'IA
16. Race After Technology
17. Student Generative Artificial Intelligence Survey 2026
18. Students are asking for AI guidance, not just policy
19. Teaching and Generative AI
20. The 2026 AI Index Report
21. The AI-Native University Must Guard Against Getting Better at the Wrong Things
22. The Impact of AI on Students' Reading, Critical Thinking, and Problem Solving
23. The Myriad Complex Ways Young People Use AI

24. The Time Constraints of AI Access Could Change How We Think
25. The Unintended Consequences of Artificial Intelligence and Education
26. Writing with machines? Reconceptualizing student work in the age of AI