GOVERNANCE WITHOUT READINESS

How 53 American Universities Responded to Generative AI

Working Paper

Evidence Snapshot: March 2026

Version 6.3

Tommy N. Turner

Williamsburg and Christiansburg, Virginia

757-784-0979

tommy@gofast.pro

linkedin.com/in/tommyturner

ORCID: 0009-0005-7422-9881

October 2026

© 2026 Tommy N. Turner

Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0). https://creativecommons.org/licenses/by/4.0/

Table of Contents

Abstract 4
I. 9.7 Percent 7
II. The Demand Signal 10
II.1 Survey convergence 10
II.2 Hiring-instrument rebuilds 11
II.3 The 83/65 gap 14
II.4 The one institutional channel that stayed open — Lehigh CCPD 14
II.5 The graduate-side evidence 15
II.6 Source tiers and methodological defense 16
III. The Institutional Response 20
III.1 The channels and their output 20
III.2 Aggregate coding across 53 institutions 21
III.3 The five-tier classification 22
III.4 Why the infrastructure sat idle at the other fifty-two 28
IV. The Readiness Gap 30
IV.1 The gap as a finding 30
IV.2 The four constituencies 31
V. The Literature: What the Field Predicted, What the Field Confirms, What the Field Will Not See 40
Organizational theory: the response was predictable 40
Cuban 1986: the present is the fifth iteration 41
Credentialing versus preparation: the incentive structure 41
Emerging GenAI governance scholarship: documented but not connected 42
The AAC&U foil cluster: the architects reading their own data 43
The Roche correspondence: the methodological limitation is the finding 45
VI. The Assessment Trap 48
VI.1 The measurement problem, not the integrity problem 48
VI.2 What the instruments were built to measure 49
VI.3 The evidence from both sides of the transom 50
VII. The Outlier 54
VII.1 Lehigh 54
VII.2 Independent-discipline convergence 55
VIII. The Disposition 59
Institutional Proximity and Scope Disclosures 59
Limitations and Future Research 60
About the Author 61

Suggested Citation 61
Companion Materials 62
Appendix A: 53-Institution Coding Table 63
Appendix B: Voices by Constituency 67
Employers and Hiring Managers 67
Faculty 68
Students 69
Career Center Professionals and Institutional Leaders 69

Abstract

American universities responded to generative AI primarily by governing it. This paper examines how 53 institutions built that response and what they did not build alongside it. Using public-source institutional profiling across a consistent ten-section framework, the study documents AI governance and workforce readiness activity at 23 nationally ranked peers of William & Mary and 30 Virginia four-year institutions, with an evidence snapshot date of March 2026.

The findings reveal a structural asymmetry. Ninety-eight percent of institutions formed AI governance committees. Ninety-two percent revised academic integrity codes. Eighty-five percent provisioned enterprise AI tools. But only 7.5 percent documented employer outreach regarding AI expectations for graduates, and 5.7 percent included career services in their AI governance structures. Of 53 institutions, one originated its AI strategy from employer input. Twenty-nine were governance-dominant: the governance channel defined their public AI posture, and their readiness work, where present, stayed almost entirely on the interior.

Drawing on Weick's loose coupling, DiMaggio and Powell's institutional isomorphism, and Cohen, March, and Olsen's garbage can model, the paper argues that existing organizational structures routed generative AI into integrity and compliance channels while excluding the labor-market signals that would have prompted a readiness response. The governance apparatus consumed the institutional bandwidth, committee time, and administrative attention that readiness would have required. Governance structures, by their success in managing risk, made the harder work of preparation invisible and untimely.

The paper identifies a deeper problem beneath the governance-readiness gap. Institutions treated AI primarily as an integrity threat, building enforcement infrastructure around the question "did the student do this work?" Generative AI has made that question insufficient. When a model can produce the artifact the credential was designed to certify, the assessment model itself requires redesign. The dominant institutional response protects an assessment framework whose validity has changed. The outlier institutions that broke the pattern share a common feature: they started from an external question about what graduates need rather than an internal question about what might go wrong.

The paper reads the outlier case at the variable where the break occurred. The single institution in the dataset that originated its AI strategy from employer input placed its career services unit at the table where the institution set AI policy. The fifty-two other career services units in the dataset have equivalent analytical capacity. What differs at the one case is procedural: the unit has standing in the channel where the institution decides. The paper also reads two independent-discipline findings that converge on the execution-gap reading from outside the governance evidence base. Correspondence with Gene Roche, a former career center director, surfaces a methodological limitation of public-source profiling that recurs as the substantive finding: career services is not reported as an institutional AI function because the field does not think of it as one. A *Journal of Macromarketing* commentary by Reza Barkhi and Michelle Seref, written from inside business school administration without reference to the governance evidence, reaches the same execution-readiness finding from a different vantage. Three readings converge on one pattern at three evidentiary distances: the cross-institutional coding, the inside informant, and the independent-discipline commentary.

Keywords: generative AI, higher education governance, workforce readiness, institutional isomorphism, assessment, academic integrity, loose coupling

JEL Classification: I23

Acknowledgments. Thank you to my wife, Charlotte of 44 years, for listening to the pitches and reading the endless drafts, and to my hound dog, Red, for his patience when I appealed to him to wait just one more minute while I finished one more sentence, as he waited by the door anxiously with his tail wagging. Thank you to Gene Roche, executive professor emeritus in the William & Mary School of Education, special advisor to the dean of the Mason School of Business on AI, and a former career center director, whose April 2026 correspondence sharpened the paper's reading of career services as a structural position and surfaced the methodological-limitation-as-finding argument in §V's Roche correspondence subsection. Remaining errors are the author's.

Section I

9.7 PERCENT

I. 9.7 Percent

In 2025 and 2026, 191 self-selected American campuses joined the AAC&U Institute on AI, Pedagogy, and the Curriculum. Each team was required to declare three to five goals for its AI action plan. The distribution they reported is the field's most authoritative public statement of what its most AI-engaged institutions believe they should be working on. Curriculum integration, 66.1 percent. Student AI literacy, 59.1 percent. Governance structures and policies, 58.1 percent. AI technology implementation, 55.9 percent. Faculty development, 50.5 percent. AI research and innovation, 48.4 percent. Responsible and ethical implementation, 42.5 percent. Student success and support services, 29.6 percent. Developing partnerships with industry, 9.7 percent. Operational efficiency, 1.6 percent. Career services, employer engagement, and workforce transition do not appear as categories.

José Antonio Bowen and C. Edward Watson published the table at page 322 of the second edition of *Teaching with AI* (Johns Hopkins University Press, 2025).[1] Immediately after the table, the authors comment that "the concern for teaching and learning is telling," and read the distribution as confirmation that the field is appropriately prioritizing pedagogy. They do not return to the 9.7 percent. The architects of the field's AI response collected the field's own data, published it in full, and called the distribution encouraging. The 9.7 percent is the one number on the page they did not interpret. This paper interprets it.

The pattern is not new. Larry Cuban's *Teachers and Machines* documented in 1986 that American education has cycled through the same exhilaration-to-disappointment sequence with film and radio in the 1920s and 1930s, instructional television in the 1950s and 1960s, and the desktop computer in the early 1980s.[2] Each cycle produced reformer hype, administrative push, slow classroom uptake, and eventual disappointment. Generative AI is the fifth iteration.

In 1922, Thomas Edison told an interviewer: "I believe that the motion picture is destined to revolutionize our educational system and that in a few years it will supplant largely, if not entirely, the use of textbooks."[3] Three years later, the schoolteacher Virginia Church answered him in verse, in a poem titled "Antiquated" published in her collection *Teachers Are People*. Her closing lines: "Education will become a matter of pressing the button. Perhaps I can get a position at the switchboard."[4]

Techno-prophet hype and the contemporary mockery of techno-prophet hype have been twinned in the American historical record for a century. The 2026 institutional response to AI belongs to that record. Section V reads Cuban and the four prior cycles in full. Section I needs only the artifacts.

---

[1] José Antonio Bowen and C. Edward Watson, Teaching with AI: A Practical Guide to a New Era of Human Learning, 2nd ed. (Baltimore: Johns Hopkins University Press, 2025), 322. Source catalog entry 53. The 9.7 percent finding for industry partnerships sits beside the authors' commentary that "the concern for teaching and learning is telling."

[2] Larry Cuban, Teachers and Machines: The Classroom Use of Technology Since 1920 (New York: Teachers College Press, 1986). Source catalog entry 26, verified primary read.

[3] Cuban, Teachers and Machines, 9, citing Henry A. Wise, Motion Pictures as an Aid in Teaching American History (New Haven: Yale University Press, 1939), 1.

[4] Cuban, Teachers and Machines, 5, citing Virginia Church, "Antiquated," in Teachers Are People (1925), via Neil Postman, The End of Education (New York: Knopf, 1996), 49–50.

This paper is built from a hand-coded dataset of 53 American universities, profiled between October 2025 and March 2026 through a consistent ten-section public-source framework with an evidence snapshot date of March 2026. The two cohorts are the 23 national peer institutions adopted by William & Mary's Board of Visitors as its benchmark cohort, and 30 Virginia four-year institutions spanning public R1 research universities, regional publics, and private colleges. The framework covers governance architecture, policy infrastructure, teaching and learning posture, tool ecosystems, academic integrity integration, research governance, library and information literacy, infrastructure and strategic investment, executive signaling, and an evidence archive of every cited source.[5]

Across the 53 institutions, 98 percent had formed AI governance committees or task forces. 92 percent had revised academic integrity codes. 85 percent had provisioned enterprise AI tools. 7.5 percent had documented employer outreach regarding AI expectations for graduates. 5.7 percent had included career services in their AI governance structures. One institution out of 53 originated its AI strategy from employer input. Twenty-nine were governance-dominant: the governance channel defined their public AI posture, and their readiness work, where present, stayed almost entirely on the interior.

The hand-coded dataset and the architects' published goals are two readings of the same field from two evidentiary angles. They report the same pattern. The architects published the 9.7 percent and called it encouraging. This paper looks at what the 9.7 percent built when it reached the institutional level, and what the other 90.3 percent built instead.

American higher education routed generative AI through the institutional channels that already existed: academic integrity, faculty governance, IT procurement, accreditation compliance. The channels routed AI to the questions those channels were built to answer. They produced governance. They did not produce readiness. The structures performed exactly as their theory of operation predicted. That is the deeper finding and the reason this paper exists.

What follows documents the demand side, the institutional response the 53 universities built, the readiness gap, the scholarly literature that predicted the response, and the assessment dimension that connects the governance failure to its root in how American universities define what a credential certifies.

[5] The dataset, ten-section framework, and evidence archive are documented in Section II below and in the William & Mary AI Inventory Report and the William & Mary Peer Benchmark v5.3 in the supporting materials.

Section II

THE DEMAND SIGNAL

II.1 Survey convergence

II.2 Hiring-instrument rebuilds

II.3 The 83/65 gap

II.4 The one institutional channel that stayed open — Lehigh CCPD

II.5 The graduate-side evidence

II.6 Source tiers and methodological defense

II. The Demand Signal

The labor market moved on AI-era hiring expectations between 2024 and 2026. The movement is documented across four independent structured-survey instruments, one institutional employer-outreach study, a body of graduate-side reporting, and the hiring practices of named executives at named companies. This section walks the evidence in order of evidentiary weight, from structured surveys to institutional outreach to individual executive testimony, and ends with the methodological accounting that names what each tier can and cannot carry.

II.1 Survey convergence

The four structured-survey instruments that read the labor market for AI-era hiring expectations between mid-2024 and early 2026 used different methods, sample frames, and commercial affiliations. They reported the same finding.

In May 2024, Microsoft and LinkedIn surveyed 31,000 workers and leaders across 31 countries for the Work Trend Index. Sixty-six percent of business leaders said they would not hire a candidate who lacked AI skills. Seventy-one percent said they would prefer a less experienced candidate with AI skills over a more experienced candidate without them. Seventy-seven percent said AI would allow early-career employees to take on greater responsibilities sooner.[6] The survey captured reported hiring preferences eighteen months after the public release of ChatGPT, from the people who make hiring decisions, in the instrument that covers the largest cross-national knowledge-worker sample available.

One year later the Microsoft follow-up described Frontier Firms — companies in which every employee, including new hires, was expected to function as an "agent boss" managing AI tools and autonomous agents as a core part of daily work.[7] The competency profile expected of entry-level employees in those firms included iteration with AI rather than acceptance of first drafts, prompting with context and intent, and judgment about when to override AI-generated output. The 2024 Work Trend Index reported the shift in hiring criteria. The 2025 follow-up reported the shift in job content.

The World Economic Forum's 2025 Future of Jobs Report surveyed over 1,000 global employers representing more than 14 million workers across 22 industry clusters and 55 economies. AI and big data ranked first among fastest-growing skill categories, with an 87 percent net increase over the five-year forecast horizon.[8] Two-thirds of employers reported plans to hire specifically for AI skills. Thirty-nine percent of core workforce skills were expected to be transformed or outdated by 2030. The forecast window opens with the graduating class of 2026.

Lightcast's analysis of AI skill requirements in job postings found that demand for generative AI skills in non-technical roles grew ninefold between 2022 and 2024. By 2024, fifty-one percent of AI-skilled job postings were outside IT and computer science, and AI skills commanded an average salary premium of

---

[6] Microsoft and LinkedIn, 2024 Work Trend Index Annual Report: AI at Work Is Here. Now Comes the Hard Part (May 2024). Survey of 31,000 knowledge workers across 31 countries conducted in February 2024.

[7] Microsoft, 2025 Work Trend Index Annual Report: Introducing Frontier Firms (April 2025).

[8] World Economic Forum, The Future of Jobs Report 2025, Insight Report (January 2025), ISBN 978-2-940631-90-2. Source catalog entry 38.

twenty-eight percent.[9] PwC's analysis of approximately one billion job advertisements found that skills in AI-exposed occupations were changing sixty-six percent faster than in other roles, up from twenty-five percent the prior year.[10] The rate of change is itself accelerating.

The four instruments differ in scale and method. The Work Trend Index is a cross-national survey of self-reported hiring preferences. The WEF report is a forward-looking forecast from a global employer sample. Lightcast and PwC are labor-market analytics built on job-posting corpora that capture what employers ask for in writing. The surveys are what employers say. The analytics are what employers type into requisition systems. They converge. II.6 names the commercial and methodological limitations of each and argues the convergence carries after the most aggressive discounting.

II.2 Hiring-instrument rebuilds

II.1 gathered what employers said. This subsection gathers what they did. Between 2024 and 2026, named executives at named companies rebuilt the mechanical procedures by which their companies screen candidates for AI-era roles. The rebuilds happened inside hiring cycles, and they are documented by the executives who built them.

Alejandro Castellano, CEO of Caddi, told *Fortune* in 2025 that interviewers at the company regularly ask candidates about their experience using AI tools and encourage candidates to use Cursor, Claude Code, or Copilot during interview exercises. His reasoning: "We want to see how they work in real conditions."[11] Caddi's interview is an instrument that permits the tool the candidate will use on the job and evaluates the interaction rather than the output.

Ehsan Mokhtari, CTO of ChargeLab, made a similar change after recognizing that candidates were avoiding AI tools during interviews on the assumption that penalties would follow. ChargeLab then formalized AI use in its hiring process and added AI literacy to departmental OKRs across support, product, sales, and operations functions.[12] The diagnosis was about the instrument rather than the candidates: the prior interview was producing false negatives because candidates were suppressing a skill ChargeLab wanted to see. The new interview reads what the old one could not.

Alan Price, Global Head of Talent Acquisition at Deel, built the most explicit instrument of the group. For non-engineering roles, Deel now evaluates AI literacy on a graduated four-level scale running from effective prompting at level one through independent tool experimentation, measurable AI-driven results, and the construction of autonomous agentic workflows at level four.[13] Price's reported assessment of the current candidate pool: "Most interviews don't get past level one." The scale reads for a skill the prior instrument did not read for.

---

[9] Lightcast 2024 report on generative AI skill demand in U.S. job postings.

[10] PwC 2025 Global AI Jobs Barometer (approx. one billion job advertisements analyzed).

[11] Alejandro Castellano, CEO of Caddi, quoted in Fortune, June 2025. Direct quote: "We want to see how they work in real conditions."

[12] Ehsan Mokhtari, CTO of ChargeLab, quoted in Fortune, June 2025. The detail about candidate self-suppression is reported in the same Fortune piece.

[13] Alan Price, Global Head of Talent Acquisition at Deel, quoted in Fortune and on LinkedIn, March 2026. Direct quote: "Most interviews don't get past level one assessment of candidates, or not test at all."

Omar Haroun, CEO of Eudia, hired a recent business graduate with no prior experience who had taught himself to use generative AI tools. The hire produced marketing-grade output at roughly ten times the rate of more experienced employees working without AI proficiency.[14] Haroun's report is an outcome observation rather than an instrument redesign. It records a hiring decision made against the grain of traditional experience-based screening and the productivity that followed.

Greg Shove of Section reported a similar observation: a single new hire used AI to self-coach in real time on tasks that peers with more than five years of experience were slower to complete.[15] The Haroun and Shove reports are individual data points, and the pattern across them is that the candidates the named executives chose to hire and watched succeed in entry-level roles were ones whose AI proficiency came to them outside institutional preparation.

Natasha Pillay-Bemath, IBM's VP of Global Talent, described the shift in what entry-level roles now require: a move from task-driven execution toward analysis, problem-solving, and responsible AI use.[16] Nolan Church, former head of talent at Google and DoorDash, named the cost of the shift directly. Companies that once helped new hires mature now expect them to "add value on Day One," with onboarding and training no longer a normal investment.[17] The instrument redesigns sit on top of a background condition Pillay-Bemath and Church together name: employers are no longer in the business of finishing what an undergraduate education started.

---

[14] Omar Haroun, CEO of Eudia, quoted in Inc., November 2025. Direct quotes from the hire profile: "zero experience," "obsessed with AI," and "10 times the output of a more experienced employee working without the boost from the new technology."

[15] Greg Shove, CEO of Section, post on LinkedIn, March 2026.

[16] Natasha Pillay-Bemath, VP Global Talent at IBM, quoted in IBM Think, March 2026.

[17] Nolan Church, former head of talent at Google and DoorDash, interview on the Riviera Partners podcast, March 2026. Direct quote: "add value on Day One."

Table 1. Hiring-instrument changes and entry-level observations reported by named executives, 2025–2026.

Name	Title	Observation	Source
Alejandro Castellano	CEO, Caddi	Interview exercises permit candidate use of Cursor, Claude Code, or Copilot; evaluation covers the interaction, not only the artifact	*Fortune*, June 2025
Ehsan Mokhtari	CTO, ChargeLab	Formalized AI use in hiring after observing candidates self-suppress; AI literacy added to departmental OKRs across all functions	*Fortune*, June 2025
Alan Price	Global Head of Talent Acquisition, Deel	Four-level AI-literacy scale for non-engineering roles: effective prompting, independent tool experimentation, measurable AI-driven results, construction of agentic workflows	*Fortune* / LinkedIn, March 2026
Omar Haroun	CEO, Eudia	Hired inexperienced AI-proficient graduate; observed ten-times output of experienced peers without AI proficiency	*Inc.*, November 2025
Greg Shove	CEO, Section	New hire used AI to self-coach in real time, outpacing peers with five-plus years of experience	LinkedIn, March 2026
Natasha Pillay-Bemath	VP Global Talent, IBM	Entry-level roles shifting from task-driven execution to analysis, problem-solving, and responsible AI use	IBM Think, March 2026
Nolan Church	Former Head of Talent, Google / DoorDash	Day-one value expectation; onboarding and training no longer a normal investment	Riviera Partners podcast, March 2026

Seven executives at seven companies. Canva (Brendan Humphreys, CTO) and Arcade (Alex Salazar, co-founder) reported comparable hiring-instrument changes in the same period; both are treated in Section VI.3, where the same evidence carries the assessment-instrument rebuild argument.

The instrument rebuild extends past the interview structure to the screening layer itself. Students report being interviewed by AI systems, not by recruiters. Platforms including CodeSignal and Eightfold analyze verbal and vocal data to predict future performance and screen applicants at scale, producing a first-pass reading no human on the hiring side now produces. The observation comes to the paper through correspondence with a former career center director who has tracked the change through student reports from his prior institution.[18] Two features of the platform shift matter for the readiness question the paper reads. First, the instrument the student first encounters in the hiring process is itself an AI system, and the competency the system is reading for is the candidate's capacity to operate inside an AI-augmented interaction. Second, the mechanism is opaque to public institutional reporting. The former director reports that his two interviews with contacts who have inside knowledge of AI

[18] Eugene Roche, personal communication to the author, April 3, April 9, and April 16, 2026. Roche is executive professor emeritus in the William & Mary School of Education, special advisor to the dean of the Mason School of Business on AI, and a former career center director. Permission to cite granted in writing April 9, 2026; re-confirmed April 16, 2026 in the author's AGB Trusteeship review request. Citation covers the April 3 observations on governance and career services, the April 9 relationship-capital mechanism and professional-staff-versus-faculty-prestige observation (the direct-quoted phrases in §III.4 are from the April 9 reply), and the April 16 observation on two off-record interviews Roche conducted with individuals holding inside knowledge of AI recruitment practices at their companies.

recruitment protocols at major employers yielded acknowledgment that the frameworks exist but declined to share operational detail. The instrument layer named in Table 1 describes hiring changes the employers report on the record. The AI-screening layer sitting upstream of those interviews is not on the record at the same resolution, and the career services units closest to the student report the change from the other side.

II.3 The 83/65 gap

Tom Watson and Lee Rainie produced two surveys for Elon University's Imagining the Digital Future Center under AAC&U co-sponsorship, one year apart. *Leading Through Disruption* (January 2025) reported a late-2024 survey of 337 senior higher education leaders.[19] *The AI Challenge* (January 2026) reported a time-series follow-up with 1,057 faculty respondents.[20]

The 2024 leadership survey asked what institutions were doing in response to generative AI. The reported distribution: 69 percent had developed AI policies, 63 percent had created task forces, 44 percent had added new AI-related courses, 19 percent had created majors or minors, and 14 percent had added AI literacy as a general education outcome. The descending order of those percentages is the empirical pattern at the leadership-decision tier — the deeper the commitment required, the fewer institutions had made it.

The same instrument asked the same leaders about their own AI use and about their graduates. 83 percent of the 337 leaders personally used generative AI tools. 65 percent of the same leaders believed their own graduates were not prepared to work in environments where AI matters.

The 83 percent and the 65 percent come from the same respondents answering the same survey on the same day. The leaders who built the institutional response to generative AI use the tools themselves at a rate of 83 percent, and do not believe, at a rate of 65 percent, that the response they built is producing graduates prepared for the work. The gap lives inside one instrument the field produced.

The 2026 follow-up expanded the respondent population to 1,057 faculty and ran the instrument against the 2024 leadership baseline. One year later, the response distribution was essentially unchanged. The graduate-readiness measures held flat. Nothing had moved.

II.4 The one institutional channel that stayed open — Lehigh CCPD

The Lehigh University Center for Career and Professional Development conducted qualitative outreach to 15 employer partners across industries in the summer of 2025. The report published from that outreach is the only direct institution-to-employer dispatch on AI expectations for graduates in the 53-institution dataset. Lehigh's CCPD is one of the four institutions that make up the 7.5 percent reported in Section I.

---

[19] Tom Watson and Lee Rainie, Leading Through Disruption: Higher Education Leaders Assess AI's Impacts on Teaching and Learning (Elon, NC: Elon University Imagining the Digital Future Center / AAC&U, January 2025), 6, 9, 12. Source catalog entry 51.

[20] Tom Watson and Lee Rainie, The AI Challenge: Higher Education Faculty Assess AI's Impacts on Teaching and Learning (Elon, NC: Elon University / AAC&U, January 2026), 13–14. Source catalog entry 52.

The report organized its findings around five themes. AI use is broad and growing across industries. AI literacy is a baseline expectation rather than a differentiator. Human judgment and communication remain critical. Ethical AI use registers as a hiring signal. AI readiness can be demonstrated without a technical background.[21] One employer told the CCPD team that new hires are expected to be using AI tools by their third day on the job.

Lehigh's provost summarized the finding in a sentence: employers want "applied fluency with AI, not model building."[22] The phrase carries the load. "Applied fluency" names a competency expected of every graduate regardless of major or intended function, and separates it from the technical-credential default career services units have tended toward when the question turned to AI: recommend a data science course, a computer science minor, a certification program. Lehigh's outreach produced a different instruction. Every graduate will need to use AI tools well in the work the graduate actually does.

The Lehigh CCPD report is a single institution's dispatch from one unit's outreach to 15 employers in one summer. It is not a structured survey, and II.6 treats the methodological weight it can and cannot carry. What it documents is the presence of one institutional channel that remained open to the labor market under AI conditions. The absence of equivalent channels across the rest of the dataset is what the paper is about.

II.5 The graduate-side evidence

II.3 reported the 83/65 gap from the administrator side. Eighty-three percent of provosts and presidents agreeing AI fluency will be required of their graduates within five years. Sixty-five percent reporting their institutions are not preparing students for it. The reciprocal evidence comes from the graduates themselves, and from the labor-market analysts who track them. Three independent instruments. One finding.

Cengage's *Graduate Employability Report* runs annually against a U.S. sample of recent four-year graduates and the employers who hire them. The 2024 instrument reported that 55 percent of graduates considered themselves unprepared for the workforce they entered.[23] The 2025 follow-on reported that 30 percent of graduates secured full-time work in their field of study, down from 41 percent the prior year.[24] The first number is self-assessment. The second is a labor-market outcome the respondent cannot spin. They move in the same direction and they move sharply.

CarringtonCrisp's *See the Future* 2026 instrument reports a parallel finding from a different vantage. Graduates rating their AI preparation, employers rating the graduates they hired, and business school deans rating their own programs converge on the same mismatch the administrator surveys describe

---

[21] Lehigh University Center for Career and Professional Development, qualitative employer outreach summary, Summer 2025. Fifteen employer partners across industries. Five thematic findings as reported in the internal CCPD report.

[22] Lehigh University provost, quoted in the CCPD Summer 2025 outreach report. Direct quote: "applied fluency with AI, not model building."

[23] Cengage Group, Graduate Employability Report 2024, finding on graduate self-reported preparedness for the workforce.

[24] Cengage Group, Graduate Employability Report 2025, year-over-year change in in-field full-time placement among recent graduates.

from the supply side.[25] The CarringtonCrisp triad does not require any one constituency's number to be clean. It requires that three constituencies who do not consult each other land on the same description.

Daniel Zhao, lead economist at Glassdoor, reframed the methodological objection to graduate self-assessment in a January 2026 *Fortune* commentary. The standard objection is that graduates complaining about their preparation are unreliable witnesses to their own readiness. Zhao's response is that the relevant question is not whether the self-report is calibrated. The relevant question is whether the labor market is treating the gap as real. By the second metric, the 2025 Cengage in-field placement collapse from 41 to 30 percent is the audit. The graduate self-reports are corroboration after the fact.[26]

Jonna Sercombe, CEO of Cegos, made the institutional version of the same point in a 2026 LinkedIn essay aimed at corporate learning leaders. Sercombe's argument is that the lack of graduate preparation for AI-augmented work is no longer in serious dispute among the employers his company serves, and that the practical question for higher education is whether any institutional channel exists to receive that signal. His answer, written from the seat of a company whose business model depends on closing the gap downstream, is that the channel does not exist at scale and that corporate learning is absorbing the cost.[27]

The graduate-side evidence and the administrator-side evidence are reciprocal readings of the same gap. The administrators report from above. The graduates report from below. Labor-market analysts read the gap sideways through placement and wage data. The employer surveys of II.1 register it as hiring demand; the executive instruments of II.2 register it as workforce strategy. Five vantages, one description. II.6 takes up the methodological objection that all five could be tilted in the same direction by the same commercial incentives, and answers it.

II.6 Source tiers and methodological defense

The demand signal in this section rests on five evidentiary tiers, each carrying a different kind of weight, each subject to different objections, none of them sufficient on its own. The case is the convergence.

The first tier is the structured employer survey. Microsoft and LinkedIn, PwC, the World Economic Forum, Cengage and CarringtonCrisp. Large samples, instrument design by research firms with reputational stakes in the result, response weighting against industry and geography. The objection is the obvious one. Each of these instruments was commissioned by an organization with a commercial interest in the answer. Microsoft sells Copilot. LinkedIn sells learning. PwC sells AI advisory engagements. Cengage sells curriculum. The instruments are not neutral, and a reader who treats them as neutral has missed something the instruments themselves do not hide.

---

[25] CarringtonCrisp, See the Future 2026, produced in association with EFMD, 4uni-solutions, and Full Fabric (February 2026). Drawing on responses from 1,863 participants across 40 countries. The study found that seven in ten alumni reported teaching themselves to use AI by downloading tools and experimenting, 64 percent learned from online videos, and 30 percent learned from friends, while a growing share of employers reported that universities were not preparing graduates with hands-on AI experience. The self-teaching finding is the piece the paper cites at the AI-fluency-acquisition-channels argument.

[26] Daniel Zhao, Fortune, 2026, on the relationship between graduate self-assessment and labor-market outcome data as audit mechanism.

[27] Jonna Sercombe, "Lack of preparation," LinkedIn, 2026. Sercombe's seat at Cegos and the corporate-learning context are load-bearing for the claim about absent institutional channels.

II.1 already conceded this and made the harder argument anyway. Discount each instrument for its commercial tilt, set the most aggressive plausible adjustment against each finding, and the surveys still converge on the same numbers. Eighty-some percent of employers reporting integration. Two-thirds reporting use as a hiring criterion. Majorities reporting that AI fluency separates candidates at the offer stage. The convergence does not depend on any single survey being clean. It depends on multiple instruments with different commercial tilts arriving at the same place after each is discounted on its own terms. That is what survey convergence is for. It is the only thing it is for.

The second tier is the hiring-instrument rebuild reported by named executives. Table 1. Seven companies, seven executives speaking on the record under their own names about specific changes their firms had made to interview structure, screening criteria, evaluation rubrics, or onboarding expectations. This tier carries weight the surveys cannot. The executives are not aggregating; they are reporting their own firms' decisions. They are not anonymous; their statements attach to professional reputations. They were not commissioned; the statements were made in trade press, conference remarks, podcast interviews, and company communications where the cost of misrepresentation is borne personally. The objection is sample size. Seven firms is not a population. The answer is that seven firms reporting structural rebuilds in the same direction within the same eighteen-month window, across industries that do not share hiring norms, is a different kind of evidence than a survey of a thousand HR managers reporting attitudes. It documents decisions, not opinions about decisions.

The third tier is the administrator survey. Watson and Rainie, *Leading Through Disruption* and *The AI Challenge*. Provosts and presidents reporting on their own institutions and on the graduates those institutions are producing. Eighty-three percent agree that AI fluency will be required of their graduates within five years. Sixty-five percent report that their institutions are not preparing students for it. The 83/65 gap is the load-bearing finding of II.3 and the load-bearing finding of the section. Its weight comes from a feature the other tiers cannot match: it is self-incrimination. The respondents are the people responsible for the gap they are reporting. The answer to "did the surveys lean the question" is that no plausible question-leaning produces a two-thirds admission of institutional failure from the people whose job it is to deny institutional failure.

The fourth tier is institutional employer outreach. Lehigh CCPD, summer 2025, fifteen employers, three weeks, one career services unit doing the work that fifty-two other institutions in the dataset did not do. II.4 treated the methodological caveats. The report is one institution's dispatch from one unit. It is not a structured survey. It documents the existence of an institutional channel, not the typical state of institutional channels. Its evidentiary work in this section is not "here is what employers want" — the surveys do that work. Its evidentiary work is "here is what it looks like when an institution actually asks." The absence of equivalent evidence across the rest of the dataset is the finding the rest of the paper is built on.

The fifth tier is graduate-side testimony. II.5 reported it. The objection is selection. The answer is that graduates reporting the same thing across institutions, industries, and graduation years, reporting it in venues where the cost of complaint is borne personally and the audience is professional rather than therapeutic, is corroboration from the constituency with the most direct exposure to the gap the surveys, the executives, and the administrators are all describing — not a self-selected sample of grievance.

No tier carries the section alone. The structured surveys can be discounted for commercial tilt. The executive statements can be dismissed for sample size. The administrator surveys can be questioned for reflexivity. The institutional outreach is one report. The graduate testimony is one constituency. Each objection is real. The case is that the same finding emerges from five sources with five different vulnerabilities, none of which overlap. Discount any one tier and the other four still converge. Discount

any two and the remaining three still converge. The demand signal is a convergence across instrument types that were not designed to corroborate each other and have no incentive to.

The institutional response that the rest of the paper documents was made against this signal. Not against an absent signal, not against a contested signal, not against a signal that arrived too late to act on. Against this one.

Section II close The signal was available. It was available in the surveys, in the executive statements, in the administrators' own self-reporting, in the one career services unit that did the work, and in the testimony of the graduates who walked into the labor market the rest of the paper is about. The fifty-three institutions catalogued in this study had access to all of it. What they did with it is the subject of Section III.

Section III

THE INSTITUTIONAL RESPONSE

III.1 The channels and their output

III.2 Aggregate coding across 53 institutions

III.3 The five-tier classification

III.4 Why the infrastructure sat idle at the other fifty-two

III. The Institutional Response

III.1 The channels and their output

The channels Section I named were institutional channels in a literal sense. They had standing. They had budgets. They had committee rosters, reporting lines, and procurement authority. When generative AI arrived as an institutional problem, these were the structures equipped to receive it, and what they produced across the 53 institutions in the dataset is visible in the public record.

The dominant output was committees. The University of Virginia stood up four of them inside the first eighteen months: a Gen-AI Task Force, a Provost's Task Force on AI Research, a Provost's Advisory Committee on AI, and school-level task forces reporting upward.[28] George Mason appointed what was reported as one of the first Chief AI Officers at a public university and assembled a visioning task force of fifty to seventy members.[29] William & Mary chartered the Artificial Intelligence Policy Initiative: twelve members across six schools, biweekly meetings for an academic year, a sixty-page report, and a recommendation for a permanent AI Innovation and Policy Council to follow.[30] Virginia Tech produced a 168-policy gap analysis against its existing policy corpus.[31] Boston College convened parallel bodies at the campus and executive levels.[32] The five committees catalog what institutional energy produces when it arrives in the committee-formation channel.

Other outputs followed the same logic. The second was policy text: syllabus templates, acceptable-use documents, honor-code amendments, data-classification frameworks. The third was enterprise tool provisioning, negotiated centrally with data-protection addenda, and the institutional vendors that resulted are the usual set: ChatGPT Edu, Microsoft Copilot, Gemini for Education. The fourth was detection-tool posture, which consumed its own institutional conversation and terminated in adoption, rejection, or deferral depending on the institution and never in consensus across the dataset.

What the channels were built to answer were questions about risk, access, and conduct. Those are the questions they answered.

---

[28] University of Virginia AI governance bodies documented in the William & Mary Peer Benchmark v5.3.2 institutional profile for UVA: Generative AI Task Force, Provost's Task Force on AI Research, Provost's Advisory Committee on AI, and school-level task forces.

[29] George Mason University. Chief AI Officer appointment and 50–70 member visioning task force documented in the Peer Benchmark v5.3.2 institutional profile for GMU.

[30] William & Mary Artificial Intelligence Policy Initiative: twelve members across six schools, biweekly meeting cadence for an academic year, sixty-page report, recommendation for a permanent AI Innovation and Policy Council. Documented in the William & Mary AI Inventory Report and the Peer Benchmark v5.3.2 W&M profile.

[31] Virginia Tech 168-policy gap analysis. VT AI Inventory v4.0.1 and Peer Benchmark v5.3.2 VT profile.

[32] Boston College parallel campus and executive-level AI committees. Peer Benchmark v5.3.2 BC profile.

III.2 Aggregate coding across 53 institutions

The coding was done from public sources: institutional websites, policy documents, board minutes, press releases, teaching center publications, accreditation reports, honor council records, and the documents produced by the AI committees themselves. Each institution was profiled across the same ten-section framework used throughout the paper and scored on four governance variables and five readiness variables. The coding captures what institutions have published. What they may be doing internally is outside its reach. A thin profile can conceal extensive unreported work. A thick profile reflects extensive public documentation and nothing stronger. What the coding reliably captures is where institutions chose to spend their public-facing AI energy.

Across the 53, fifty-two institutions had formed a named AI committee, task force, advisory group, or working group with documented membership or charge. Forty-nine had revised an honor code, academic integrity policy, or conduct code with AI-specific language. Forty-one had extended a formal data classification or data governance framework to AI tool usage. Forty-five had provisioned an enterprise AI tool with institutional data protection. Forty had built a structured faculty development program for AI integration in teaching, and not only one-off workshops.

Four of the 53 had documented outreach to employers specifically about AI expectations for graduates. Three had a career services unit with a documented role in institutional AI governance or programming. One originated an AI strategy from employer input.[33]

Three of the four governance variables sit above seventy percent of the dataset; the fourth, the syllabus template, sits at 37.7 percent. The two external-facing readiness variables sit below eight percent. The interior readiness variables track the governance side of the ledger: new programs, faculty development, and enterprise provisioning are additions to what the institution already teaches and already runs, scored off the institution's own existing inventory of activity. The external-facing readiness variables ask something different. They ask whether the institution has gone outside its own walls to the employers its graduates are about to work for and adjusted in response to what those employers said they needed. Four institutions did. Three connected the question to career services at the governance level.

Governance-dominant, in the classification used here, is a reading of the whole institutional profile rather than a count. It names the institution where the governance channel defines the public AI posture and readiness work, where present, stays on the interior. Twenty-nine of the 53 institutions sit there at the March 2026 evidence snapshot. The indicators bound the classification without determining it. Twenty-six of the twenty-nine carry three or four of the four governance indicators. Twenty-five carry at least one readiness indicator, most often enterprise provisioning and faculty development, and one, Tufts, carries a career-services role in AI strategy. A strict count test (three or more governance indicators, at most one readiness indicator, nothing external-facing) is met by seven institutions, all in the Virginia stratum. Another twelve built readiness programming further on the interior variables while the governance channel continued to define the institution's public AI posture. The remaining twelve distribute across the other three tiers.

---

[33] Counts across 53 institutions from the 53-institution coding table in supporting materials, Summary Counts tab. Governance variables: AI committee or task force 52/53 (98.1%); integrity code revised 49/53 (92.5%); data classification framework 41/53 (77.4%); syllabus template 20/53 (37.7%). Readiness variables: enterprise AI tool provisioned 45/53 (84.9%); faculty development program 40/53 (75.5%); new AI degree or certificate 25/53 (47.2%); employer outreach documented 4/53 (7.5%); career services in AI strategy 3/53 (5.7%). Variable definitions and coding rules on the Coding Notes tab of the same workbook.

Table 2. Summary of governance and readiness coding across 53 institutions, March 2026.

Variable	Count	Percent
Governance variables		
AI committee or task force	52 / 53	98.1%
Integrity code revised for AI	49 / 53	92.5%
Data classification framework	41 / 53	77.4%
Syllabus template	20 / 53	37.7%
Readiness variables		
Enterprise AI tool provisioned	45 / 53	84.9%
Faculty development program	40 / 53	75.5%
New AI degree or certificate	25 / 53	47.2%
Employer outreach documented	4 / 53	7.5%
Career services in AI strategy	3 / 53	5.7%
Tier distribution		
Tier 1 (employer-originated)	1 / 53	1.9%
Tier 2 (substantial readiness + governance)	7 / 53	13.2%
Tier 3 (readiness present, governance dominates)	12 / 53	22.6%
Tier 4 (governance-dominant)	29 / 53	54.7%
Tier 5 (minimal documented activity)	4 / 53	7.5%

The two external-facing readiness variables — employer outreach documented and career services in AI strategy — sit at 7.5 percent and 5.7 percent against a governance-side range of 37.7 to 98.1 percent. The 9.7 percent §I named from the architects' own declared-goals table is the closest structural match the literature provides for the external-facing column on the field's published data. The 9.7 and the 7.5 / 5.7 are reading the same pattern from two evidentiary angles.

III.3 The five-tier classification

The classification the coding produces is five tiers. What separates one tier from another is the ratio between the governance infrastructure the institution built and the readiness infrastructure it built, read across the whole profile and informed by the nine variables named above rather than computed from them. Every indicator can be checked against the source register in the companion dataset. The tier lines are the author's reading of the full profiles, and the descriptions below name the institutions whose indicators sit outside their tier's usual pattern.

Tier 1 is the institution whose AI strategy originated from employer input. One institution in the dataset met the test. Lehigh University's Center for Career and Professional Development conducted fifteen structured employer interviews, and the resulting Lehigh AI Fluency Framework was built to what the employers described.[34]

[34] Lehigh University Center for Career and Professional Development, "Employer Expectations for AI Fluency," summer 2025 (supporting materials). Fifteen structured employer interviews conducted during spring and summer 2025 as empirical ground for the Lehigh AI Fluency Framework. See also the Lehigh institutional profile in the Peer Benchmark v5.3.2 and Lehigh AI Inventory v0.9 (both in supporting materials). This is the only instance in the 53-institution dataset of an AI strategy originating from

Tier 2 is the institution with substantial readiness infrastructure running alongside the governance infrastructure. Seven institutions sit here. All seven documented all three interior readiness variables, and six of the seven scored high on the governance variables; Marymount documented two of the four. Two, Clemson and George Mason, also reached the external-facing side. The governance channel is present at Tier 2 at the same scale as at Tier 4. What is different at Tier 2 is the readiness layer sitting beside it.

Tier 3 is the institution with readiness programs present but governance still dominating the institutional AI posture. Twelve institutions sit here. All twelve built faculty development programs and provisioned enterprise tools, and ten also created new AI-specific programs or certificates. One, Washington & Lee, documented employer outreach on AI; the other eleven did not cross to the external-facing side of the ledger. The governance channel is the loudest institutional AI voice at Tier 3 even where the readiness work is non-trivial.

Tier 4 is the institution where the governance channel is the institutional AI voice and readiness work, where it appears on the public record, does not define the institution's AI posture. Twenty-nine institutions sit here. Tier 4 is the modal cell in the dataset. The institutions here are not inactive. They are doing the work the channels through which they routed AI were built to do.

Tier 5 is the institution with minimal documented AI activity across all nine variables. Four institutions sit here. These are small colleges and under-resourced publics where a full-scale institutional response did not occur on the public record during the evidence window.

The distribution is one, seven, twelve, twenty-nine, four. Tier 4 is both the median and the modal cell. The institution the dataset points to when asked what an American university built in response to generative AI is an institution at Tier 4. III.3 walks the tiers in reverse order, Tier 5 to Tier 1, with institutional profiles.

III.3.1 Tier 5

Four institutions sit at Tier 5: Emory & Henry University, Hollins University, Sweet Briar College, and UVA's College at Wise. The public record at each is thin across all nine coding variables. Emory & Henry is the only institution in the dataset whose coding carries no formed AI committee. It also carries no AI-specific integrity-code revision on the public record and no documented detection-tool position. The other three have a committee on record and not much beyond it. These are small private colleges and an under-resourced public branch campus where the institutional capacity that produced 98 percent committee formation at the rest of the dataset did not generate comparable public-facing output during the evidence window. The thin profile is a function of scale and resources. Tier 5 is named here because the classification is exhaustive. The finding of §III does not live in Tier 5 and the tier is not walked further.

III.3.2 Tier 4

Twenty-nine institutions sit at Tier 4 at the March 2026 evidence snapshot. That is too many to profile individually. Five cases walk the range.

---

documented employer input. The interview protocol was a five-question framework (Q1 entry-level AI skills and tools; Q2 differentiating competencies and mindsets; Q3 emerging tools, platforms, and use cases by industry; Q4 power skills for leveraging AI; Q5 additional preparation gaps), with five themed findings consolidated in the Summary of Findings section.

Brandeis published a single issue of Brandeis Magazine during the 2024–2025 academic year that carried two of its own professors speaking about AI in their classrooms. Elanah Uretsky, Associate Professor of International and Global Studies, said she did not use AI in her classes and did not permit her students to use it on assignments, citing a conviction that the capacity to think without AI assistance was a civic matter.[35] Aldo Musacchio, Professor of Business at the Brandeis International Business School, in the same issue of the same magazine, described using AI to speed prototyping work in his entrepreneurship classes.[36] Both professors operated under the same Brandeis institutional AI guidance at the same moment in the same academic year. The governance layer at Brandeis was doing everything the public record at a Tier 4 institution shows it was doing. The student experience at the course level across those two classrooms was incoherent in a way the governance layer had no mechanism to resolve. The delegation mechanism is the mechanism.

James Madison University states the dynamic on the record. Provost Heather Coltman sent an email to faculty in August 2023 directing each instructor to formulate an individual syllabus AI policy and acknowledging in the same email that students taking multiple courses in the same semester "might receive contradictory guidelines."[37] The acknowledgment was candid and accurate. The institution recognized the incoherence the delegation mechanism would produce and delegated anyway. At Brandeis the dynamic is documented by external journalism. At JMU it is documented by the provost.

Boston College sits at Tier 4 and the parallel-committees configuration named in §III.1 is the reason: BC built the governance layer at both the campus and the executive levels and did not build external-facing readiness during the evidence window.

Hampden-Sydney College mandates verbatim AI policy language in every course syllabus and does not permit instructors to edit the language.[38] The small-college case makes the Tier 4 argument in a different register from the elite national-peer case. The governance posture at Hampden-Sydney is maximal. The instrument is a required template enforced by administration with no instructor variance, and the readiness side of the coding shows a faculty development program and nothing more. What Hampden-Sydney shares with Brandeis is what places both at Tier 4: a governance channel that sets the public AI posture, and readiness work that stops short of a new credential and of anything external-facing.

---

[35] Elanah Uretsky and Aldo Musacchio, faculty quotes on AI use in teaching, Brandeis Magazine, Winter 2024/2025 issue, AI-quotient feature at https://www.brandeis.edu/magazine/2025/winter/featured-stories/ai-quotient.html. Single footnote covers both Brandeis quotes because both appeared in the same issue of the same publication during the same academic year.

[36] See note 34.

[37] Heather Coltman, Provost, James Madison University, email to faculty dated August 15, 2023, directing individual instructors to formulate course-level AI policy and acknowledging that students "might receive contradictory guidelines." Referenced in The Breeze and WHSV coverage. Peer Benchmark v5.3.2 and Virginia Tech AI Inventory v4.0.1 carry cross-institutional syllabus-template context.

[38] Hampden-Sydney College mandatory syllabus template with verbatim AI policy language, instructor edits not permitted, updated June 23, 2025. Liberty University "The Liberty Way" student-code amendment adding AI-specific conduct provisions and authorizing detection-tool enforcement. Both documented in the Virginia institution profiles assembled for the 30-institution Virginia cohort. Consolidated footnote because both institutions are Tier 4 small-institution cases illustrating different governance-posture variants in the same paragraph.

Liberty University amended its student code to add AI-specific conduct provisions and authorized detection tools as part of the enforcement infrastructure.[39] The governance posture at Liberty is explicitly restrictive rather than delegative. Tier 4 contains institutions whose governance channel defaults to instructor discretion and institutions whose governance channel defaults to prohibition. The variable the coding tracks is whether the external-facing readiness layer is present, and at both kinds of Tier 4 institutions it is absent.

The five cases are representative and not exhaustive. A list of twenty-nine institutions would not illustrate the cell better than five cases walked. Twenty-nine is the finding. The institution an American college student is most likely to attend during the March 2026 evidence window is an institution where the governance channel built the AI infrastructure that defines its public posture, and the readiness channel, where it built anything, built it on the interior.

III.3.3 Tier 3

Twelve institutions sit at Tier 3. These built readiness programming on top of their governance work and, with one exception, did not cross to the external-facing side of the readiness ledger. The exception is Washington & Lee, which documented employer outreach on AI. No Tier 3 institution carried a career center with a documented role in institutional AI strategy at the March 2026 evidence snapshot.

Brown University built faculty development through its Sheridan Center for Teaching and Learning and provisioned enterprise AI tools through institutional procurement. Duke University launched faculty programming through its Learning Innovation and Lifetime Education office and created AI-specific curricular programming through its Trinity College and Pratt School of Engineering. Emory University built readiness programming through its Center for Faculty Development and Excellence and deployed enterprise AI capacity through the Office of Information Technology.[40] These three national peers share the same Tier 3 profile: a teaching-and-learning center running the readiness side, an IT or information-services organization running enterprise provisioning, and a governance layer that continued to define the institution's public AI posture as the readiness side expanded.

Georgetown, the University of Southern California, and American University follow the same logic with institution-specific variants. Georgetown built faculty development through its Center for New Designs in Learning and Scholarship. USC built programming through its Center for Generative AI and Society. American University's Center for Teaching, Research and Learning carried faculty development alongside the Kogod School of Business's AI curricular launches.[41] Each of these institutions has a governance-layer public-facing posture that absorbed more institutional AI attention during the evidence window than the readiness-layer work did.

---

[39] See note 38.

[40] Tier 3 institutional profiles in supporting materials (Peer Benchmark v5.3.2) for Brown University (Sheridan Center for Teaching and Learning), Duke University (Learning Innovation and Lifetime Education; Trinity College and Pratt School of Engineering AI curricular programming), Emory University (Center for Faculty Development and Excellence; Office of Information Technology enterprise AI provisioning), Georgetown University (Center for New Designs in Learning and Scholarship), University of Southern California (Center for Generative AI and Society), and American University (Center for Teaching, Research and Learning; Kogod School of Business AI curriculum). Consolidated footnote.

[41] See note 40.

Washington & Lee University is the small-college case at Tier 3: faculty development built through the Harte Center for Teaching and Learning, enterprise AI provisioning, and AI-specific curricular programming on a scale appropriate to the institution's size. Washington & Lee also documented employer outreach on AI, the one Tier 3 institution to do so, while the governance channel still set its public AI posture. The tier assignment tracks the variable ratio and not the absolute magnitude of the programming.

The twelve-institution cell also includes Virginia Commonwealth University, Stony Brook University, and three institutions already named in §III.1 for their committee-formation output: the University of Virginia, Virginia Tech, and William & Mary. Each of the three §III.1 institutions sits at Tier 3 on the coding. The readiness-side work at each (AIPI's policy infrastructure at William & Mary, Virginia Tech's 168-policy gap analysis handed off to its Pamplin College of Business AI programming, UVA's school-level task-force outputs and Darden Institute programming) is non-trivial and is documented at length in the Peer Benchmark profiles. None of the three crossed to external-facing readiness during the evidence window.

Tier 3 is the cell where the governance-versus-readiness gap narrows without closing. The governance channel remained the loudest institutional AI voice at Tier 3 even where the readiness work was substantial.

III.3.4 Tier 2

Seven institutions sit at Tier 2. Each documented all three interior readiness variables alongside its governance work, and two, Clemson and George Mason, also reached the external-facing side. The seven distribute into two clusters that reached Tier 2 by structurally different paths.

The supply-side cluster built AI capacity through institutional investment in new programs, new schools, or existing AI infrastructure. Rice University launched a Bachelor of Science in Artificial Intelligence in Fall 2025. Vanderbilt University received a $25 million gift to establish a College of Connected Computing with a dedicated founding dean and deployed the Amplify AI platform across the institution. The University of Georgia operates an Institute for Artificial Intelligence that predates most of the contemporary AI landscape, launched an undergraduate AI Certificate open to all majors, and runs an AI for Business certificate through the Terry College.[42] Northeastern University built AI programming at scale through Khoury College of Computer Sciences and expanded the college's AI-specific curricular infrastructure.[43] The four supply-side institutions reached Tier 2 through the same structural path: institutional resources converted into new AI-capable programs that face outward toward the labor market those programs feed.

The demand-side cluster is smaller and the path is different. Clemson University built its TEACH-AI program as a cross-college faculty development initiative that included a career services representative

---

[42] Tier 2 supply-side cluster programming documented in supporting materials (Peer Benchmark v5.3.2): Rice University Bachelor of Science in Artificial Intelligence launched Fall 2025; Vanderbilt University $25 million gift establishing the College of Connected Computing with founding dean and Amplify AI platform deployment; University of Georgia Institute for Artificial Intelligence (established 1984), undergraduate AI Certificate, Terry College of Business AI for Business certificate; Northeastern University Khoury College of Computer Sciences AI programming.

[43] See note 42.

from the Center for Career and Professional Development alongside faculty and librarians.[44] The Career Innovation Lab that followed gave students structured access to generative AI tools for career preparation. Clemson's governance-layer work was substantial and present alongside the readiness work, which is why Clemson is at Tier 2 rather than Tier 1. What puts Clemson on the readiness side of the ledger at all is that its career services unit had a seat at the governance table when strategy was being set. Marymount University launched Virginia's first Bachelor of Science in Artificial Intelligence in Summer 2024. An Applied AI Graduate Certificate followed. Marymount also secured an NSF grant specifically for faculty development in AI teaching across disciplines.[45] Marymount is a smaller Virginia private without the endowment scale of the national research universities in the supply-side cluster. The programs were built anyway.

The seventh Tier 2 institution is George Mason University, named in §III.1 for its Chief AI Officer appointment and 50–70 member visioning task force. GMU's readiness-side work includes AI program development at the College of Engineering and Computing and its Institute for Digital InnovAtion programming. The institution sits at Tier 2 on the coding and its role in §III.1 as a committee-formation illustration does not exhaust its profile.

The supply-side path to Tier 2 requires institutional resources of a kind most institutions in the dataset do not have. The demand-side path requires a career services unit with a seat at the governance table and a willingness to let what employers said reach the curriculum. The structural cost of the demand-side path is lower than the structural cost of the supply-side path. The governance-dominant majority did not take either path.

III.3.5 Tier 1

Lehigh University is the only institution in the 53-institution dataset whose AI strategy originated from documented employer input. The mechanism is the Lehigh Center for Career and Professional Development and what it did during the summer of 2025.

Fifteen employers from a cross-section of industries sat for structured qualitative interviews conducted by CCPD staff over a three-week period.[46] The interviews produced five themed findings: AI use is broad and growing across industries, AI literacy is a baseline expectation rather than a bonus, human judgment and communication remain critical, ethical use of AI is a hiring signal, and AI readiness can be demonstrated by students outside of technical majors. CCPD organized the findings and brought them to the provost's office. The provost used them to frame Lehigh's institutional AI response.

The AI Fluency Framework took its curricular shape from the employer findings. Faculty AI development programming built out from the same source. The three-tier AI Readiness program structure followed. The College of Business designated the 2025–2026 academic year its "Year of

---

[44] Clemson University TEACH-AI program membership, including career services representation from the Clemson Center for Career and Professional Development alongside faculty and librarians. Career Innovation Lab student programming. Peer Benchmark v5.3.2 Clemson profile.

[45] Marymount University Bachelor of Science in Artificial Intelligence launched Summer 2024 as the first BS in AI at any Virginia institution; Applied AI Graduate Certificate; NSF grant titled "Accelerating AI Workforce Development: Preparing Faculty for Teaching AI Across Domains and Pathways." Marymount AI Inventory v2.3 and supporting materials (AI Higher Ed Reference Guide v1.7).

[46] See note 21.

AI@Work." Each component traces its origin to what the employers described to CCPD in the interview transcripts, not to the governance-committee outputs that populated the same academic year at Lehigh's peer institutions.

Lehigh also opted out of Turnitin's AI detection product for the 2025–2026 academic year on the grounds that the false-positive rate compromised student-teacher trust.[47] The detection-tool opt-out is a governance-layer decision at every other institution that made it. At Lehigh it travels with the employer-originated strategy as an expression of the same institutional posture. An institution that built its AI strategy from what employers said they needed did not reach for a detection tool as the instrument that answered whether students were using AI.

Every institution in the 53-institution dataset has a career center. Every one of them has employer relationships with companies that hire their graduates. The infrastructure CCPD used at Lehigh in the summer of 2025 exists at the other fifty-two. It was used once.

III.4 Why the infrastructure sat idle at the other fifty-two

Gene Roche, executive professor emeritus in the William & Mary School of Education, special advisor to the dean of the Mason School of Business on AI, and a former career center director, named the mechanism in correspondence in April 2026.[48] Employers treat their interviewing and selection techniques as proprietary. A placement-side practitioner who develops trusted relationships with recruiters over years accumulates the kind of information that could shape career center policy and curricular policy, and the condition of keeping the access is discretion about what gets said publicly. Information shared once in confidence can, in Roche's words, "dry up very quickly" if mishandled. The silence at the other fifty-two institutions records a relationship structure that produces its value through private channels and loses its value when those channels become public. The silence is a condition of the work, held by the practitioners whose work depends on it. Roche confirmed the pattern empirically from outside the public record: he interviewed two individuals with inside knowledge of AI recruitment practices at their companies, and both acknowledged proprietary AI frameworks for interviewing and evaluating students while declining to share details. The proprietary layer exists. The public-source methodology this paper runs cannot see it.

The 5.7 percent figure §III.2 documented measures how many institutions place career services in a position where its labor-market information enters the governance channel as strategy. It does not measure how many career services units have that information. At fifty-two of fifty-three institutions, the practitioners closest to the labor-market signal hold information the governance channel cannot reach and that public methodology cannot see. The hierarchy the institution runs on carries the rest. As Roche observes, the real governance question is whether senior leaders understand the contribution of "professional staff who have deep knowledge but are not afforded the prestige of 'faculty.'"[49]

---

[47] Lehigh University opt-out of Turnitin AI detection for the 2025–2026 academic year, citing false-positive rates and the concern that detection tools compromised student-teacher trust. Lehigh AI Inventory v0.9 and Peer Benchmark v5.3.2 Lehigh profile.

[48] See note 18.

[49] See note 18.

Section IV

THE READINESS GAP

IV.1 The gap as a finding

IV.2 The four constituencies

IV. The Readiness Gap

IV.1 The gap as a finding

The 53 institutions in the dataset divide across the governance and readiness channels at a ratio of roughly thirteen to one. Twenty-nine were governance-dominant, and at twenty-eight of them nothing on the public record reached the external-facing side of readiness. Twelve more ran governance and interior readiness programming alongside each other, eleven of them with nothing on the external-facing side. Four had documented employer outreach specifically about AI expectations for graduates. One originated its AI strategy from employer input. §III read the distribution as a five-tier taxonomy. §IV reads the same distribution as a claim about what American higher education defaulted to when generative AI arrived as an institutional problem. The ratio reflects a sectoral pattern. The 53 institutions were selected as a cohort precisely because they are visible, well-resourced, and self-reporting at a level that makes public documentation possible. If the ratio among the 53 is thirteen to one, the ratio across the field is the same or worse.

Governance and readiness are answers to two different questions. Governance is about conduct. It specifies what the institution permits and the consequences of violation. It decides the procurement, integrity, and classification questions that novel tools raise. Readiness is about capability. It specifies what the graduate should be able to do at the point of hire, measured against the conditions of the working environment that hire will place them in. Governance is answered inside the institution about the institution's own processes. Readiness is answered across the institution's boundary to employers and back. §III showed that 98 percent of institutions answered the governance question and 7.5 percent answered the readiness question. This paper is about which question the institution defaulted to when it could only staff one.

The mechanism §III documented is a channel mechanism. Governance is the channel the institution already had. The integrity office had an address, the faculty senate had a meeting schedule, IT procurement had standing vendor relationships, and the accreditation liaison knew which quarterly reporting deadline the new guidance needed to hit. Each of those structures was equipped to receive a novel institutional problem, and the institution routed generative AI through them. Readiness would have required a new channel. The career services office at the typical American university has no standing seat at the AI governance table. It does not hold procurement authority. Its integration with the academic side of the house rarely extends beyond the minimum that placement reporting requires. Building readiness infrastructure would have meant building that channel first and running the AI response through it second. The institution defaulted to the channel that was already paved.

The default is the trap. The governance channel is the infrastructure the institution already has. The readiness channel would require building something new. Under conditions of institutional uncertainty, energy goes to the channel that already exists. That is a structural fact about how American higher education allocates bandwidth, and it is the same fact Cuban documented in the four prior technology-adoption cycles. What makes the mechanism a trap is not that the institution did the wrong work. It is that the work the institution had infrastructure to do was not the work the graduates needed done. The trap closes without a decision getting made. No one in the institution chooses to skip readiness. The decision is already inside the architecture of which channel exists and which would have to be built, and the channel that exists wins by default before the question reaches a meeting agenda.

The four constituencies §IV walks are students, faculty, employers, and career services staff. Each sits at a different position around the gap §III counted. §III showed the gap from the coding table. The

subsections that follow show it from the working lives the gap routes through. What changes across §IV.2 is the vantage.

IV.2 The four constituencies

A constituency, in the sense this section uses the term, is a group whose working life passes through the governance-to-readiness gap §III documented and whose experience of the gap is structurally different from the other groups'. The four constituencies §IV walks are students, faculty, employers, and career services staff. Two are inside the institution. Two are at the institutional boundary, one outside looking in and one inside looking out. The four subsections that follow read the gap from each position in turn.

§III classified institutions against the nine governance-and-readiness variables. §IV.2 classifies groups of people against the single variable of their structural position relative to the gap. The two taxonomies are perpendicular. An institution sits in one tier, but the four constituencies sit inside every tier, and the experience of each constituency differs by tier in ways §IV.2.1 through §IV.2.4 will name where the difference is load-bearing.

The walks are not exhaustive. Administrators, trustees, parents, accreditors, and the publishing research community each experience the gap in ways that could support analyses of their own. The four §IV walks are the four whose experience is closest to the primary evidentiary base of the paper: the 53 institutional profiles, the demand-side surveys §II documented, the Lehigh CCPD employer-interview record, and the AAC&U goal distribution §I reads. §IV.2.1 through §IV.2.4 walk them in the order named.

IV.2.1 Students

At a Tier 4 institution, the modal cell of the dataset, a student encounters the institutional AI response as a revised honor code, a published syllabus paragraph on permitted use, an enterprise-tool account accessed through single sign-on, and a set of classroom rules that change from one course to the next depending on which instructor is teaching. Twenty-nine of the 53 institutions sit in that cell. The student experience at those institutions is the modal student experience the dataset documents. The response the institution built is the response the student receives.

The governance layer defines what the student is permitted to do with AI and the consequences of violation. It does not define what the student should be able to do with AI at graduation. The governance answer is a conduct answer. The student question is a capability question, and the two answers do not substitute for each other. A student whose AI exposure across four undergraduate years is thin at a governance-dominant institution graduates knowing the institution's rules and without the practice-based fluency the labor market is starting to treat as the baseline for AI-augmented professional work. The difference between a student who has read about AI and a student who has used AI iteratively for thirty hours a week across a semester is the difference between information and practice, and the difference shows up at the point of hire.

AI fluency is a practice skill. Polanyi called it tacit knowledge: the kind of knowing that sits inside doing and cannot be fully transferred by description.[50] The name applies to working with generative AI without modification. A student can be told what prompting is and what iterative multi-pass workflow

[50] Michael Polanyi, The Tacit Dimension (Chicago: University of Chicago Press, 1966). Polanyi's original formulation that "we can know more than we can tell" anchors the tacit-knowledge tradition this paper draws on for the AI-fluency-as-practice-skill argument.

looks like without being able to do either when a task lands on the desk. The gap between knowing about and knowing how is a gap time and repetition close. Students acquire AI fluency through four channels: coursework that requires AI practice, co-curricular engagement such as research assistantships and competitions, self-directed practice outside the curriculum, and labor-market experience after graduation.[51] The institution controls the first channel at scale. Coursework is the only channel a university can design and require across an entire undergraduate cohort, and Tier 4 institutions used it to publish rules rather than to build practice-based fluency.

A student who leaves an undergraduate program with thin practice-based AI exposure enters the labor market against peers whose exposure was thick, and the gap widens after graduation rather than closing, because the thick-exposure cohort accumulates the kind of experience the thin-exposure cohort has not yet started. The tacit-craft argument names the mechanism the degradation curve: a shortfall that deepens across the first years of working life as the labor market assimilates AI-augmented practice as the baseline for junior work. The compounding is the part the institution cannot intervene in after graduation. The undergraduate years are the intervention window the institution owned.

The institutional scope of the shortfall is visible on the same surface §I reads. Nine point seven percent of the 191 AAC&U Institute campuses named industry partnerships as a goal in their AI action plans. The other 90.3 percent did not. Ninety out of every hundred students enrolled at an AAC&U Institute campus were enrolled at an institution whose AI action plan did not include a commitment to reach the employers whose hiring decisions would condition the students' first years after graduation. Read from the student side, the 9.7 percent is the share of students whose institutions made the readiness question a public commitment of their AI response. The 7.5 percent §III documented is the share of students whose institutions actually did the outreach work. The second number is smaller than the first. The student side of the gap is the difference between being enrolled at an institution that named the question and being enrolled at an institution that did not.

IV.2.2 Faculty

The governance channel §III.1 described did its work through faculty. Honor-code revisions passed through faculty senates, syllabus language through individual instructors at the course level, academic-integrity cases through committee adjudication and desk-level paper review, and enterprise-tool and detection-tool decisions through the course staff at the classrooms where the tools would and would not be permitted. Each was an answer to a governance question the institution framed at the policy level and then asked the executing layer to implement. The institution gave the governance answer. The faculty executed it.

A faculty voice catalog assembled from public coverage between January 2024 and March 2026 documents six distinct categories of response to the institutional AI posture: Refusal, Reluctant Compliance, Cautious Experimentation, Active Integration, Evangelism, and Governance Burden.[52] The

---

[51] Internal synthesis supporting the tacit-craft argument. The four-channel student-fluency-acquisition frame (coursework, co-curricular engagement, self-directed practice, labor-market experience) and the degradation curve are consolidated from prior research notes.

[52] Four Constituencies — Faculty Voice Catalog, compiled from public coverage January 2024–March 2026. The six categories are Refusal and Resistance, Reluctant Compliance, Cautious Experimentation, Active Integration, Evangelism, and Governance Burden. Named exemplars for each category are documented in the catalog file; this paper uses the typology as a framing device and cites individual voices separately as they appear in the text.

catalog is a typology of the stances faculty have taken publicly. The stances are unevenly distributed, and the middle of the typology is where most faculty at most institutions sit. The modal response at the modal institution is Reluctant Compliance or Cautious Experimentation. Refusal and Active Integration are the tails. Evangelism and Governance Burden are smaller still, and they are the categories where the institutional-level strain is most visible. Governance Burden, in particular, names a specific experience: faculty serving on AI task forces, benchmarking against peer institutions, reviewing the national landscape, and drafting recommendations as a process that consumes significant uncompensated labor across disciplines.

The Governance Burden category has a concrete public record. The University of South Carolina's 2024–2025 Provost's Task Force on AI reported that reviewing the national landscape of AI use, benchmarking against peer institutions, and drafting implementation recommendations consumed substantial faculty labor across disciplines over the task force's term.[53] Bryan Alexander, speaking on the *Teaching in Higher Ed* podcast in January 2026 about the state of AI inside postsecondary education, described the question of figuring out what people are doing with AI on campus as a great challenge, because faculty surveyed about their AI use have all kinds of incentives to not respond accurately.[54] Both voices name the same operational reality. The governance channel consumed faculty time at the institutional-inventory level and at the self-reporting-distortion level, and the time was not compensable.

The "instructor decides" policy posture is where the governance channel and the delegation converge. When a governance-dominant institution adopts that posture without also building faculty development at scale and a mechanism for curricular coherence, it transfers to individual faculty a decision only the institution is positioned to make. The instructor-decides reading presents, on its face, as a preservation of faculty autonomy over pedagogy. The reading holds only if the scaffolding is present. At a Tier 4 institution, the scaffolding is partial at best. Twenty of the twenty-nine run a faculty development program; the coherence half of the scaffolding is left to the individual instructor. The decision the individual faculty member makes about AI in the classroom sits in the absence of the support the decision would need to produce a coherent graduate outcome across the curriculum. A syllabus that permits AI in one course and prohibits it in the next is not incoherent at the course level; each instructor is making a reasoned decision inside the constraints of the class being taught. The incoherence is at the graduate level, where the student who took both courses is walking out into the labor market with the sum of two decisions that did not talk to each other. Coherence at the graduate level was an institutional variable the institution chose not to set.

Gene Roche, executive professor emeritus in the William & Mary School of Education, special advisor to the dean of the Mason School of Business on AI, and a former career center director, reframes the

---

[53] Report of the 2024–2025 Provost's Task Force on Artificial Intelligence, University of South Carolina, October 6, 2025, https://sc.edu/about/offices_and_divisions/cte/teaching_resources/docs/ai_task_force_report.pdf. The task force's process description (reviewing the national landscape of AI use, benchmarking against peer institutions, and recommending implementation steps) is cited here from the faculty voice catalog entry under the Governance Burden category.

[54] Bryan Alexander, interview on Teaching in Higher Ed podcast, episode "Peak Higher Ed: AI's Possible Futures with Bryan Alexander," host Bonni Stachowiak, aired January 8, 2026, https://teachinginhighered.com/podcast/peak-higher-ed-ais-possible-futures-with-bryan-alexander/. Quoted language ("a really great challenge because if you polled people, they have all kinds of great incentives to not respond accurately") from the faculty voice catalog entry under the Governance Burden category, paraphrased in the §IV.2.2 text.

incoherence finding from the inside of the faculty layer.[55] Institutional governance, Roche observes in correspondence in April 2026, has rarely been the mechanism through which authentic student learning is improved. The unit of change that matters is the individual faculty member in the disciplinary setting, and the gap between a provost's AI policy framework and what happens in a classroom is wide in most institutions and especially wide across disciplinary contexts. The empirical grounding Roche cites is Daniel F. Chambliss and Christopher G. Takacs's *How College Works*, a multi-year longitudinal case study of one small liberal arts college in which the unit of analysis is the individual student and the finding is that the faculty member in the disciplinary setting is the highest-leverage variable in what the student actually learns.[56] The Chambliss finding, applied to the AI response, reads the delegation §IV.2.2 documents at a different level. The delegation is an unfunded transfer of institutional work to the faculty layer and a routing of the AI response through the layer the scholarship identifies as the one that matters, in a form the institution has declined to resource. What public-source institutional profiling undercounts at the faculty-in-disciplinary-setting level is the level at which readiness either forms or does not. The governance apparatus operates one level above that, and the distance between the two levels is the distance the AI response has to travel before it reaches the student.

The delegation §III documented arrived as unpaid work. Across the 53 institutions, faculty updated syllabus language, revised assessment approaches, evaluated enterprise tools, decided detection-tool posture, responded to student questions about permitted use, and navigated academic-integrity ramifications of ambiguous cases. Each item in that list is a standing responsibility the faculty member carried alongside the teaching and research obligations that existed before generative AI arrived. At almost no institution in the dataset was this work compensated as additional labor, and at no institution was it treated as a temporary transition cost bounded in time. The governance answer produced real work for the executing layer. The institution did not fund the work.

The institution routed the AI problem to the faculty. The routing was not itself a choice the faculty made. The decision about which channel received the AI response was made upstream, in the logic §IV.1 named: the governance channel existed, the readiness channel would have required building, and the institution defaulted to what existed. The faculty working inside a governance-dominant institution are holding the downstream consequences of that default every day they teach. The governance channel the institution built is the channel the faculty now maintain as a standing layer of the job, and the readiness channel the institution did not build is the absence the faculty encounter at every classroom point where a student asks what AI use the graduate labor market will expect of them. The delegation is the experience of the gap from the inside of the executing layer.

Inside-pair close

§IV.2.1 and §IV.2.2 read the gap from the two positions inside the institution. The student experiences it as the response the institution built and the response the institution declined to build, absorbed across four undergraduate years. The faculty experience it as the delegation the institution routed through them, absorbed across committee work and individual classroom judgment the institution declined to make centrally. The two experiences are coupled. The faculty decisions about AI in the classroom are the

---

[55] See note 18.

[56] Daniel F. Chambliss and Christopher G. Takacs, How College Works (Cambridge, MA: Harvard University Press, 2014). Multi-year longitudinal case study at Hamilton College using the individual student as the unit of analysis. The finding that faculty members in the disciplinary setting are the highest-leverage variable in what students learn is the empirical grounding Roche cites for the last-mile observation in the April 9 correspondence.

decisions the students encounter as the institutional AI posture, and the student fluency that results is the fluency the faculty were positioned to build under whatever support the institution gave them. §IV.2.3 and §IV.2.4 move to the institutional boundary: the position of employers outside and the position of career services inside-looking-out.

IV.2.3 Employers

Employers do not see the committees the institution built. Employers see the graduates. The institutional AI response arrives at the employer side of the gap as whatever the graduate can or cannot do on the first Monday of the first job. §II documented the employer signal in survey form: hiring managers across sectors weight AI fluency heavily and describe the gap between training and expectations as wide. §IV.2.3 reads that signal through named voices who run hiring and describe what they observe from the employer side of the institution-to-labor-market boundary.

AI-augmented professional work is an iterative multi-pass workflow with a human in the middle.[57] The worker prompts a model, reads the output, evaluates where it is right and where it is wrong, revises, and iterates until the result meets the standard the worker is trying to produce. The first draft arrives in seconds. The rest of the cycle is judgment. Senior work has become making judgments about AI output across rounds of iteration, and the junior work that used to teach how the craft functions is the work the machine now does. The work has changed, and the employer is reading the change.

Jennifer Spire, partner and CEO of Preston Spire, a Minneapolis advertising agency, described the inflection point at the staffing-plan level in a September 2025 interview with *MarTech*.[58] Her agency's expanding use of AI had been eliminating work that used to teach junior marketers how marketing functions, and she had been treating the elimination as a productivity gain. She then looked at her 2029 and 2030 staffing plans and saw a pipeline problem the productivity gain had made invisible. She described asking where the mid-level and higher individuals in any discipline were going to come from and concluding her agency was heading toward a crisis. Preston Spire now weights strategic thinking and AI fluency beyond basic prompt writing in its junior hires, and Spire is meeting with University of Minnesota marketing faculty to push for curriculum change. The detail that completes the observation is that Spire reports other agency CEOs at a national consortium had not yet thought about the pipeline question. The inflection point is not widely recognized inside the employer side of the boundary either.

Nikki Slowinski, EVP of Talent Experience and Development at Publicis, launched a program called Ignite at Publicis Digital Experience to cover foundational professional skills the company decided it could not wait for new graduates to bring.[59] The program has run at least five cohorts. Paul Phillips at

---

[57] Internal synthesis supporting the operational definition of AI-augmented professional work. The single-pass versus multi-pass workflow distinction, the human-in-the-middle iteration pattern, and the process-infrastructure framing consumed in §IV.2.3 are consolidated from prior research notes.

[58] Constantine von Hoffman, "AI threatens entry-level marketing jobs and the future talent pipeline," MarTech, September 8, 2025, https://martech.org/ai-threatens-entry-level-marketing-jobs-and-the-future-talent-pipeline/. Jennifer Spire interviewed as partner and CEO of Preston Spire. The 2029/2030 staffing-plan inflection point, the hiring-criteria change, the University of Minnesota marketing-faculty meetings, and the national-CEO-consortium detail are all reported in the interview and attributed to Spire. Quoted language paraphrased in the §IV.2.3 text.

[59] Three named employer voices cited together at the build-not-buy convergence point. (a) Nikki Slowinski, EVP of Talent Experience and Development at Publicis, speaking at the From Day One conference (July 2025), reported in "If AI Starts Replacing Entry-Level Jobs, How Will Future Leaders Get Started?," From

Avanade runs academies to retrain talent from non-traditional backgrounds because the pipeline for the skills Avanade needs is not reliable, and he is converting roughly 80 percent of interns to full-time hires through the pipeline he built in the absence of a reliable external one. Eric Trickett described Dropbox's GenAI Skills Academy as the posture of a company that had recognized it could not hire its way out of an AI skills gap because the skills were too new to be reliably present in the hiring pool. Across three employers and three internal programs, the posture is the same. The employer absorbs the transition the institution declined to absorb.

Trent Cotton, Head of Talent Acquisition Insights at iCIMS, observed in a June 2025 workforce report that the Class of 2025 often brings sharper AI instincts than tenured hires and warned recruiters against defaulting to shortcuts that screen those candidates out.[60] The observation cuts against a uniform crisis reading and is worth naming on its own. The employer-side finding is that the gap is real and the variance is large. A given graduate may arrive with practice or only with the rules about practice the governance channel distributed, and the employer has no way to know before the hire. The unpredictability is itself a finding about what the institution produced. The institution produced a cohort whose fluency depends on which courses a given student happened to take.

The employer observes the gap from outside the institution looking in. The observation is routed through the hire and routed back through the internal program the employer builds, and no part of the observation reaches the institutional channel where the institution makes its AI decisions. §II documented the signal in survey form. §IV.2.3 names it in the voices of the people who hire. What the employer sees is the output. What the institution built is not visible from there.

IV.2.4 Career Services

Career services staff hear the labor market before anyone else in the institution and cast no vote in the institution's AI decisions. The director of employer relations sits in the meetings where recruiters describe what they are screening for. The career coach sits with the graduating senior who asks what to put on a resume and whether the accounting role the recruiter mentioned is still an accounting role. The position has direct line of sight to the gap §III counted and no standing at the governance table §III.1 described. The combination is structural.[61]

Day One, September 17, 2025, https://www.fromdayone.com/stories/2025/9/17/if-ai-starts-replacing-entrylevel-jobs-how-will-future-leaders-get-started. (b) Paul Phillips, Global Head of HR, Talent Acquisition, and Onboarding at Avanade, speaking at the From Day One conference (May 2025), reported in the same From Day One article. (c) Eric Trickett, VP of Global Talent Acquisition and Workforce Planning at Dropbox, speaking at the CodeSignal Beyond 2025 conference (November 2025), reported in "Highlights from Beyond 2025," CodeSignal blog, https://codesignal.com/blog/highlights-from-beyond-2025/. Each of the three describes an internal employer program built because external hiring could not reliably supply the required skills. Cited together as convergence, not as a spectrum.

[60] Trent Cotton, Head of Talent Acquisition Insights at iCIMS, "iCIMS Insights June Workforce Report," press release, June 24, 2025, https://www.icims.com/company/newsroom/juneinsights2025/. Cotton's observation that the Class of 2025 may bring sharper AI instincts than tenured hires is cited in §IV.2.3 paragraph 5 as the counter-voice to a uniform crisis reading.

[61] Internal synthesis supporting the career-services-as-policy-takers frame. The policy-taker-not-policy-maker argument, the line-of-sight-without-standing claim, and the downstream-signal framing consumed in §IV.2.4 are consolidated from prior research notes.

§III documented the quantitative form of the combination. Three of the 53 institutions had career services named in their AI strategy. The 5.7 percent is the lowest value any variable in the coding table takes. At 94.3 percent of institutions, the constituency closest to the labor-market signal sits in the policy-taker position, while the governance channel §III.1 described holds the policy-maker position and does not include career services. The descriptive name for the structural position is downstream. The signal travels from employer to recruiter to career services, and the channel where the signal would need to land to change the curriculum is upstream from career services and not connected to it.

Line of sight without standing produces a specific experience of the gap. Career services staff see what the labor market is asking for through the recruiting conversations they have every week, and they do not see the governance channel, because the governance channel does not run through their office. They are positioned to know what employers want from graduates before any other constituency inside the institution knows it, and positioned to do the least about it. The experience is consistent across every institution in the dataset where the structural position holds. It is not consistent with the institutional AI decision-making process.

Bob Beaudet, Director of Employer Relations at the Peter J. Tobin College of Business at St. John's University, described in a 2025 institutional interview what employers are telling him about AI expectations: employers expect some baseline of AI tool knowledge from graduates even when they do not specify which tools, and some employers have internal AI platforms graduates are expected to navigate from day one.[62] Beaudet also described softening in accounting and financial roles as employers wait to see how AI reshapes the work. The signal Beaudet is relaying is one step removed from the hiring manager. The institution is receiving that signal through a relay, and the relay sits in a position that is not part of the institution's AI governance structure. The signal is available. The institution is not set up to hear it.

The Lehigh CCPD record is the single institutional case in the dataset where career services initiated the AI response. §III.3.1 documented that Lehigh's Center for Career and Professional Development conducted fifteen structured employer interviews in spring and summer 2025, and the resulting Lehigh AI Fluency Framework was built to what the employers described. The counterexample works because CCPD had the standing to route the employer signal upstream. At one of 53 institutions, the inside-looking-out position was connected to the governance channel. At the other 52, the line of sight existed and the connection to the channel where decisions get made did not. What Lehigh had that the other 52 did not was a seat at the decision table for the career services office, and the seat is what turned the line of sight into an institutional AI strategy.

The career services office sits in two concurrent streams of information the institution's AI committee does not receive. One is the recruiter stream: direct reports of what employers are screening candidates for. The other is the student stream: questions about what to study and what to put on the resume. The career services office absorbs both streams and does not sit in the meetings where the institution decides what to do with them. The chair career services occupies holds the information and no vote.

---

[62] "AI's Impact on Students Entering the Job Market," Johnnies Blog, St. John's University, 2025, https://www.stjohns.edu/news-media/johnnies-blog/ai-impact-students-entering-job-market. Bob Beaudet, Director of Employer Relations at the Peter J. Tobin College of Business, interviewed for the institutional blog. Quoted language paraphrased in the §IV.2.4 text.

Boundary-pair close

§IV.2.3 and §IV.2.4 read the gap from the two positions at the institutional boundary. Employers see it from outside looking in, as what the graduate cannot do on the first Monday of the job. Career services sees it from inside looking out, as what employers tell recruiters they expect and students arrive unable to deliver. The two observations describe the same gap from opposite sides of the boundary the institution routes its AI response across, and neither observation reaches the channel where the institution makes its AI decisions. §IV returns to the gap §IV.1 named. The four walks are a single account of what a thirteen-to-one ratio looks like from inside the working lives it routes through.

Section close

The four walks describe a single gap from four positions around it, and the vantage determines how clearly the gap is seen. The career services office sees it first and most directly, because the labor market signal lands there before anyone else in the institution hears it. The governance channel sees it least, because the channel was built to answer a conduct question and a capability gap does not register as a problem inside a conduct frame. The student and the faculty sit inside the institutional AI response itself, the student as its recipient and the faculty as the layer the institution routed the response through, and experience the gap as a condition of the working life rather than as something to identify against an external standard. The thirteen-to-one ratio §IV.1 named is the quantitative form of the asymmetry. The four walks are its qualitative form. The ratio holds across every vantage, and the experience of the gap tilts toward the constituencies positioned to do the least about it. §V reads the finding against the literatures the field has built: organizational theory on how institutions route novel problems through the channels they already have, higher-education scholarship on credentialing and workforce preparation, and the historical record of prior technology-adoption cycles the current one repeats.

Section V

THE LITERATURE: WHAT THE FIELD PREDICTED, WHAT THE FIELD CONFIRMS, WHAT THE FIELD WILL NOT SEE

Organizational theory: the response was predictable

Cuban 1986: the present is the fifth iteration

Credentialing versus preparation: the incentive structure

Emerging GenAI governance scholarship: documented but not connected

The AAC&U foil cluster: the architects reading their own data

The Roche correspondence: the methodological limitation is the finding

V. The Literature: What the Field Predicted, What the Field Confirms, What the Field Will Not See

The literature on this paper's question separates into three domains and one foil. The three domains — organizational theory on disruptive technology adoption, higher education scholarship on the credentialing-versus-preparation tension, and the emerging body of generative AI governance studies — explain why the pattern documented in Sections II and III was predictable from the moment ChatGPT cleared a million users. The foil is a five-source cluster published by AAC&U, Elon University, and their affiliated authors between January 2025 and early 2026. The cluster is the field's most authoritative reading of its own data. It reaches the opposite conclusion this paper reaches from substantially the same evidence base.[63]

Organizational theory: the response was predictable

Karl Weick's 1976 analysis of educational organizations as loosely coupled systems is the foundational text.[64] Weick describes institutions in which structural elements (departments, classrooms, administrative offices) are connected but retain substantial independence. The implication for the AI response is direct. A provost can issue a memo. A governance committee can approve a syllabus template. Curriculum redesign cannot move that way. Curriculum requires concurrence across faculty, departments, deans, and accreditors, and the loose coupling that lets central policy propagate as a signal is the same loose coupling that prevents central policy from propagating as integrated pedagogical change. Governance is what the central apparatus can produce alone. Readiness requires coordination the structure resists.

DiMaggio and Powell's 1983 paper on institutional isomorphism explains the convergence of those governance responses across the dataset.[65] Three mechanisms operate simultaneously. Coercive isomorphism appears in Virginia in the form of Governor Youngkin's Executive Order 30, in compliance demands from accreditors, and in the federal data privacy regime. Mimetic isomorphism is the most visible mechanism in the data. Institutions facing radical uncertainty about a new technology copy the governance artifacts of their peers, which is why three-tier syllabus templates appear at institution after institution without central coordination. Normative isomorphism flows through professional networks (EDUCAUSE, the teaching center directors, the provostial conference circuit), and these networks transmit governance practice efficiently because governance is what they were built to share. Meyer and Rowan's companion paper supplies the structural conclusion: organizations adopt formal structures to conform to legitimacy expectations in their environment, and those structures may be decoupled from what the organization actually does.[66]

---

[63] Methodological note. This section's organization into three domains and one foil follows the source catalog and AAC&U foil cluster analysis developed for this paper.

[64] Karl E. Weick, "Educational Organizations as Loosely Coupled Systems," Administrative Science Quarterly 21, no. 1 (1976): 1–19.

[65] Paul J. DiMaggio and Walter W. Powell, "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields," American Sociological Review 48, no. 2 (1983): 147–160.

[66] John W. Meyer and Brian Rowan, "Institutionalized Organizations: Formal Structure as Myth and Ceremony," American Journal of Sociology 83, no. 2 (1977): 340–363.

Cohen, March, and Olsen's garbage can model names the deeper mechanism.[67] In their formulation, universities are organized anarchies in which problems, solutions, participants, and choice opportunities flow independently and connect when they happen to collide. Generative AI collided with a set of existing institutional anxieties — academic integrity, faculty authority, data security — that already had infrastructure built around them. AI was routed through the solutions that were already running when it arrived. The integrity apparatus did not have to be invented. It was waiting.

Bromley and Powell on the contemporary expansion of compliance,[68] and Hannan and Freeman on structural inertia,[69] complete the explanation. Honor code revisions, AI policy drafting, detection tool evaluation, and committee formation consume the same institutional attention readiness work would have required. Compliance and adaptation compete for the same scarce bandwidth, and compliance has the older claim. The 53 institutions in the dataset are not failing to respond to AI. They are responding in precisely the way the theory predicts.

Cuban 1986: the present is the fifth iteration

Larry Cuban's *Teachers and Machines* documents the structural pattern of technology non-adoption across four successive classroom technologies: film and radio in the 1920s and 1930s, instructional television in the 1950s and 1960s, and the first generation of desktop computers in the early 1980s.[70] The pattern is the same in each case. Reformer hype, administrative push, slow and limited classroom adoption, eventual disappointment, and a four-stage cycle Cuban names directly: exhilaration, scientific credibility, disappointment, and teacher-bashing.[71] The Edison and Church artifacts that open this paper come from chapter 1 of Cuban's volume, where they bracket his account of the film cycle. At page 3, Cuban writes that in the insistent quest for productivity and efficiency, "the lecture, film, radio, television, and microcomputer are first cousins."[72] Generative AI is the fifth cousin. Any reviewer claiming AI is uniquely disruptive has to engage Cuban first.

Credentialing versus preparation: the incentive structure

Slaughter and Rhoades on academic capitalism document how universities have come to treat the credential as a revenue-generating asset that must be protected.[73] AI governance protects the credential. Workforce preparation serves employer interests whose returns are diffuse and delayed and harder to

---

[67] Michael D. Cohen, James G. March, and Johan P. Olsen, "A Garbage Can Model of Organizational Choice," Administrative Science Quarterly 17, no. 1 (1972): 1–25.

[68] Patricia Bromley and Walter W. Powell, "From Smoke and Mirrors to Walking the Talk: Decoupling in the Contemporary World," Academy of Management Annals 6, no. 1 (2012): 483–530.

[69] Michael T. Hannan and John Freeman, "Structural Inertia and Organizational Change," American Sociological Review 49, no. 2 (1984): 149–164.

[70] Cuban, Teachers and Machines, 5–6. The four-stage cycle is in Cuban's own page-cited prose.

[71] See note 70.

[72] Cuban, Teachers and Machines, 3, citing Raymond E. Callahan, Education and the Cult of Efficiency (Chicago: University of Chicago Press, 1962), chs. 4–7.

[73] Sheila Slaughter and Gary Rhoades, Academic Capitalism and the New Economy: Markets, State, and Higher Education (Baltimore: Johns Hopkins University Press, 2004).

capture in institutional budgets. Labaree's analysis of American education's competing goals sharpens the same point.[74] Of the three goals he identifies (democratic equality, social efficiency, social mobility), social mobility has come to dominate, which means the degree functions primarily as a positional credential whose value lies in its scarcity and signaling function rather than in the skills it conveys. Randall Collins on the credential society[75] and Bryan Caplan on signaling[76] push the argument to its empirical limits. If the degree's primary function is signaling, the institutional response of protecting the signal through integrity governance is rational behavior, and the failure to invest in workforce capability is a feature of the system rather than a malfunction of it.

Accreditation is the structural reinforcement. Eaton's overview of US accreditation[77] and Kelchen on higher education accountability[78] document the regime: input metrics, process compliance, integrity safeguards, and enrollment and completion data. Workforce readiness for an AI-integrated economy is not a category in the accreditation apparatus. Institutions invest where they are measured. Joseph Aoun's *Robot-Proof*, revised in 2024 to address generative AI, proposes a "humanics" curriculum integrating technological, data, and human literacies.[79] The proposal is the most fully developed counter-vision in the literature. Its isolation as a single institutional leader's framework is itself the diagnostic.

Emerging GenAI governance scholarship: documented but not connected

The third domain is sparse and skewed. The peer-reviewed work that exists is mostly comparative and document-analytic, and it documents what policies exist rather than whether the policies produce workforce-ready graduates. Jin et al.'s 2025 comparative analysis of 40 universities across six regions finds universal emphasis on academic integrity, ethical guidelines, and assessment redesign for misuse mitigation.[80] Parker et al.'s 2025 analysis of 343 leading universities across five Anglophone countries finds wide variation in governance approaches but a uniform emphasis on integrity over employer-aligned competency building.[81] An et al.'s 2025 analysis of the top 50 US universities documents that 94 percent of institutions delegated AI policy-setting to course-level faculty discretion: the loose-coupling

---

[74] David F. Labaree, "Public Goods, Private Goods: The American Struggle over Educational Goals," American Educational Research Journal 34, no. 1 (1997): 39–81.

[75] Randall Collins, The Credential Society: An Historical Sociology of Education and Stratification (New York: Academic Press, 1979; reissued New York: Columbia University Press, 2019).

[76] Bryan Caplan, The Case Against Education: Why the Education System Is a Waste of Time and Money (Princeton: Princeton University Press, 2018).

[77] Judith S. Eaton, "An Overview of U.S. Accreditation," Council for Higher Education Accreditation, revised 2020, https://www.chea.org/overview-us-accreditation.

[78] Robert Kelchen, Higher Education Accountability (Baltimore: Johns Hopkins University Press, 2018).

[79] Joseph E. Aoun, Robot-Proof: Higher Education in the Age of Artificial Intelligence, rev. ed. (Cambridge, MA: MIT Press, 2024).

[80] Yueqiao Jin et al., "Generative AI in Higher Education: A Global Perspective of Institutional Adoption Policies and Guidelines," Computers and Education: Artificial Intelligence (2025).

[81] L. Parker et al., "Comparative Analysis of Artificial Intelligence Policies in Higher Education Institutions," Information Retrieval Journal (Springer, 2025), DOI: 10.1007/s10791-025-09745-5.

finding restated as quantitative data.[82] Wu, Zhang, and Carroll on Big Ten case studies document multi-unit governance architectures.[83] Luo's critique of the originality assumption underlying GenAI integrity policies is among the few sources that explicitly connects governance design to workforce relevance.[84]

The EDUCAUSE 2024 AI Landscape Study, the most authoritative inside-the-institution survey, documents that policy formation has substantially outpaced curricular integration.[85] The 2026 EDUCAUSE follow-up reports that workforce skills development at most institutions means upskilling existing employees rather than transforming student preparation.[86] Across the entire emerging literature, the connection between institutional AI governance and graduate workforce outcomes is documented in survey form by industry-affiliated organizations (Cengage Group, NACE)[87] and asserted in trade-press commentary, but it has not been the subject of large-scale peer-reviewed empirical work. That is the gap this paper occupies.

The AAC&U foil cluster: the architects reading their own data

A five-source cluster published by AAC&U, Elon University, and their affiliated authors sits across the literature as foil rather than support. The 9.7 percent finding that opens this paper sits inside this cluster. Section I exhibits the number. This section places it in the architect-tier context that turns it from anomaly into pattern.

Watson and Rainie's *Leading Through Disruption* surveyed 337 senior higher education leaders in late 2024.[88] The structural-response distribution at page 9 is the empirical pattern in a bar chart: 69 percent of institutions had developed AI policies, 63 percent had created task forces, 44 percent had added new AI-related courses, 19 percent had created majors or minors, and 14 percent had added AI literacy as a general education outcome. The descending order is the argument. The institutional-operations list in the same survey instrument names recruitment, financial matters, student life, athletics, fundraising, and alumni engagement. Career services is not on the list. The instrument did not ask. Page 12 reports that 83 percent of leaders personally use generative AI tools while 65 percent believe their own graduates were not prepared to work where AI matters. The leaders integrated the tools into their own workflow and built the governance apparatus to restrict everyone else. Watson and Rainie's *The AI Challenge*

---

[82] Yun-Jo An, Jae Hwan Yu, and Schuyler James, "Investigating the Higher Education Institutions' Guidelines and Policies Regarding the Use of Generative AI in Teaching, Learning, Research, and Administration," International Journal of Educational Technology in Higher Education 22, no. 10 (2025), DOI: 10.1186/s41239-025-00507-3.

[83] Chuhao Wu, He Zhang, and John M. Carroll, "AI Governance in Higher Education: Case Studies of Guidance at Big Ten Universities," Future Internet 16, no. 10 (2024): 354.

[84] Jieyu Luo, "A Critical Review of GenAI Policies in Higher Education Assessment: A Call to Reconsider the 'Originality' of Students' Work," Assessment & Evaluation in Higher Education 49, no. 5 (2024): 651–664.

[85] Jenay Robert, 2024 EDUCAUSE AI Landscape Study (Louisville, CO: EDUCAUSE, 2024).

[86] Jenay Robert, The Impact of AI on Work in Higher Education (Louisville, CO: EDUCAUSE, 2026).

[87] Cengage Group, 2025 Graduate Employability Report (Boston: Cengage Group, September 2025); National Association of Colleges and Employers, Job Outlook 2025 (Bethlehem, PA: NACE, 2024).

[88] See note 19.

(January 2026) is the time-series follow-up across 1,057 faculty.[89] One year later, the structural response distribution is essentially unchanged, the graduate-readiness measures held flat, and the career services exclusion repeats in the follow-up instrument. Same omission, one year later, same authors, same template.

Bowen and Watson's *Teaching with AI*, second edition, reports at page 322 the declared-goals distribution from 191 self-selected AI-engaged American campuses participating in the 2025–2026 AAC&U Institute on AI, Pedagogy, and the Curriculum.[90]

Curriculum integration	**66.1%**
Student AI literacy	**59.1%**
Governance structures and policies	**58.1%**
AI technology implementation	**55.9%**
Faculty development	**50.5%**
AI research and innovation	**48.4%**
Responsible and ethical implementation	**42.5%**
Student success and support services	**29.6%**
Developing partnerships with industry	**9.7%**
Operational efficiency	**1.6%**

Career services, employer engagement, and workforce transition do not appear as categories. Bowen and Watson, immediately after publishing the table, comment that "the concern for teaching and learning is telling," and they read the distribution as confirmation that the field is appropriately prioritizing pedagogy. They do not return to the 9.7 percent. As Section I notes, this is the one number on the page they did not interpret. Same data, opposite verdict, same paragraph, same page, written by one of the field's most cited authors. Full-text search of the volume's text returns zero hits for "career services," "career counseling," "experiential learning," "co-op," "workforce development," or "employer engagement." The authors see the labor market clearly. They name the AI skills salary premium, the 66 percent of business leaders who say they would not hire without AI skills, and the entry-level labor market collapse. They locate the entire response in the classroom.

Shyr's "How do you do AI?" feature in *Liberal Education* surveys eight faculty voices across Caltech, Smith, Kennesaw State, Holyoke Community College, AUBG, MIT, and UMass Amherst.[91] Career services is absent across roughly 3,000 words. The Caltech framing is diagnostic. Frederick Eberhardt, professor of philosophy and co-director of Caltech's Linde Center for Science, Society, and Policy, frames the faculty debate about AI as a classroom-permission question. Caltech. The institution that trains the people building the systems frames AI as an evaluation problem. Anderson, Watson, Rainie, and Anderson's *Student Guide to Artificial Intelligence: AI-U 2025* devotes a two-page spread at pages 16–17 to a "career plan for the age of AI" that quantifies World Economic Forum workforce stakes and then lists the recommended actions: cultivate mentors, take classes, seek leadership in student organizations, join group research, prioritize internships, complete certification programs, build a portfolio, join student

[89] See note 20.

[90] See note 1.

[91] Daphne Shyr, "How Do You Do AI?" Liberal Education (Fall 2025). Catalog entry 54.

AI groups.[92] Every action is curricular, individual, or student-club activity. Not one action involves career services as an institutional function. The document is endorsed by NASPA, EDUCAUSE, ALA, the Online Learning Consortium, and AMICAL. The professional associations whose own members are accountable for career services outcomes endorsed a document that hands the workforce transition problem to the eighteen-year-old reader as a personal development project.

The foil cluster is the inside-leg evidence at national scale, in the field's own voice, from the field's own architects. The 53-institution dataset that anchors this paper documents the same pattern in a hand-coded sample. Both evidence tiers are required.

The Roche correspondence: the methodological limitation is the finding

Gene Roche, Executive Professor Emeritus in the William & Mary School of Education and a former career center director, read the paper in early April 2026 and made two observations about the public-source methodology.[93] Both initially read as limitations of the inside-leg evidence. Both turn into findings.

The first is the last-mile point. Institutional governance, Roche notes, has rarely been the mechanism through which authentic student learning is improved. The unit of change is the individual faculty member in the disciplinary setting. The outlier institutions documented in Section VII may have better governance posture, but the readiness they are building still has to travel that last mile, and faculty-level integration work is largely invisible to public-source institutional profiling. The implication is that the framework may be undercounting readiness activity at the level where it matters most. The implication is also that the governance apparatus the paper documents is operating at a level above the level where teaching actually changes, which sharpens rather than weakens the central argument.

The second is the career services point. Roche reports running multiple employer-engagement initiatives at W&M's career center in his prior role, none of which "would show up in public documents, for many complex cultural reasons." Other schools, he suggests, may be in the same position. If the public record at W&M cannot see Roche's own employer-engagement work, the public record at the other 52 institutions cannot see analogous work either. The reasonable inference is that the methodology systematically undercounts career services activity. The harder inference, which the AAC&U cluster supports independently, is that the cultural invisibility itself is the finding. Watson and Rainie's institutional-operations list excluded career services in 2024 and again in 2025. Bowen and Watson's monograph contains zero references to career services across an 11,000-line text extraction. Anderson et al.'s Student Guide treats career services as a function students must perform on their own behalf. The methodological limitation Roche identifies recurs as a structural exclusion in three independent sources at the architect tier. The field does not report on career services as an institutional AI function because the field does not think of career services as an institutional AI function. The methodological limitation is the substantive finding.

The literature predicts the pattern. The historical record establishes that the present-day institutional response is the fifth iteration of a documented cycle. The credentialing scholarship explains the incentive structure. The emerging GenAI governance literature documents the policy adoption without yet linking it to outcomes. The AAC&U foil cluster shows the field's most authoritative architects reading the same

---

[92] Janna Quitney Anderson, Tom Watson, Lee Rainie, and Hope Anderson, Student Guide to Artificial Intelligence: AI-U 2025, 2nd ed. (Elon, NC: Elon University / AAC&U, 2025), 16–17. Catalog entry 55.

[93] See note 18.

data and reaching the opposite verdict. The Roche correspondence converts what looked like a methodological limitation of the inside-leg evidence into a confirming finding from the architect tier. What the literature does not yet contain is empirical work bridging institutional AI governance and workforce preparation across a structured American sample. That is the space this paper occupies.

Section VI

THE ASSESSMENT TRAP

VI.1 The measurement problem, not the integrity problem

VI.2 What the instruments were built to measure

VI.3 The evidence from both sides of the transom

VI. The Assessment Trap

VI.1 The measurement problem, not the integrity problem

Academic integrity was already running when ChatGPT arrived. It had a staff, a policy manual, a detection vendor, and a standing committee. It was built to answer one question: did the student do this work. The governance response routed AI through that question because that question had a home.

The question has an answer in principle. It no longer has an answer in practice. Detection is unreliable at the individual artifact level, and the human costs of unreliable detection are on the record.[94] But even a clean detection result does not resolve the underlying problem, because the underlying problem is not who produced the artifact.

The essay, the problem set, the take-home exam, the term paper, the portfolio, the lab report, the coding assignment — each was built as a measurement instrument. Each was understood as a reliable proxy for a cognitive operation the curriculum claimed to develop: argument, synthesis, recall, computation, judgment. The grade aggregated the readings. The transcript aggregated the grades. The degree aggregated the transcript. Every layer of the credentialing stack rests on the construct validity of the instruments at the bottom.

Generative AI can produce most of those artifacts at a quality that passes the reading. Not all of them. Enough of them. The grammar is clean, the argument is coherent, the citations are plausible, the calculations are right when the math is routine. The artifact is still legible. The artifact is no longer a reading of the operation, because the operation can be skipped.

This is construct-validity failure at the artifact layer. The instrument has lost its correlation with the trait. A student who wrote every word by hand in 2026 is still producing an artifact whose validity as a measurement reading is unknown, because the instrument's capacity to distinguish the trait from its absence was what gave the reading its meaning, and that capacity is gone. The student did not change. The instrument did.

Governance cannot rebuild the instrument. The assignment, the rubric, the exam, and the grade live inside the course, with the instructor. The governance layer approves syllabus templates, publishes policies, contracts with detection vendors, and revises honor codes. None of those actions touches the instrument. What the governance layer can act on is who is allowed to produce the artifact and under what conditions. Integrity names a problem the governance layer owns the levers to act on. Measurement names a problem the governance layer cannot reach.

---

[94] Debora Weber-Wulff et al., "Testing of Detection Tools for AI-Generated Text," International Journal for Educational Integrity 19 (2023), Article 26, doi: 10.1007/s40979-023-00146-z. The controlled test of fourteen commercial AI-text detectors found accuracy at or below 50 percent in distinguishing AI-generated from human-written text under adversarial conditions, and concluded that the tools are neither accurate nor reliable at the individual-artifact level. Yale University's Poorvu Center for Teaching and Learning has published guidance to the same effect, noting that Yale does not support use of AI-detection tools because no method detects AI use with certainty. Lehigh University's 2025–2026 opt-out of Turnitin AI detection, documented at §III.3.5 and §VII.1, is consistent with the reliability finding and with the broader teaching-center position. Detection unreliability at the individual-artifact level is the factual claim §VI.1 paragraph 2 rests on.

The two readings produce opposite responses. Integrity produces enforcement: detect, restrict, punish, preserve. Measurement produces redesign: new instruments, new rubrics, new aggregation, and eventually a new account of what the degree certifies. The data from 53 institutions shows overwhelming investment in the first response and almost none in the second. That distribution is what the governance layer is equipped to produce when the question is measurement and the only levers it owns are integrity levers.

VI.2 What the instruments were built to measure

Each of the seven instruments named in §VI.1 was built to measure a cognitive operation a student could perform alone. The mapping was not uniform: the essay carried argument most heavily, the problem set carried computation, and the other instruments sat at various weights of the remaining operations. No syllabus stated this because no syllabus had to. The instruments were the operational definition of the operations, and the operations were the operational definition of what the curriculum claimed to develop.

All five operations are pre-AI capabilities. Argument, synthesis, recall, computation, and judgment were the vocabulary of human intellectual work in an environment where a student sitting alone with a blank page, a textbook, a laboratory bench, or a code editor was the unit of production. The instruments were calibrated against that unit. A reading produced inside the instrument counted as a reading of the student because no other entity in the room could have produced it. The floor of what a credentialed adult could be expected to do alone was the floor the instruments read against.

The capability question and the conduct question §IV.1 named run through each instrument in the inventory. Every instrument measures a specific cognitive capability. Whether the student produced the artifact alone is a separate question, and it is a conduct question. The instrument was designed to read the capability on the assumption that the student produced the artifact alone, because in the environment the instrument was built for there was no other way the artifact could exist. The conduct question was not a question the instrument had to answer. When the operation runs, the instrument measures the capability. When the operation is skipped, the instrument records whatever the skipping process produced. The instrument cannot tell the two cases apart. It was built in an environment where the two cases were the same case.

The employer signal §II named reads the same five operations against a different floor. Microsoft's Work Trend Index reported 66 percent of business leaders saying they would not hire a candidate who lacked AI skills and 71 percent saying they would prefer a less experienced candidate with AI skills over a more experienced candidate without them. The 2025 Frontier Firms follow-up named the specific competencies expected of entry-level employees: iteration with AI rather than acceptance of first drafts, prompting with context and intent, and judgment about when to override AI-generated output. Lightcast documented a ninefold increase in generative AI skill demand in non-technical job postings between 2022 and 2024 and a 28 percent salary premium on the postings that demanded the skill. The floor the labor market treats as the operational baseline is the same five operations performed in iteration with AI as a working partner, where judgment now applies to AI output the student did not produce. Multi-pass human-in-the-middle iteration is the operational vocabulary for the work the labor market hires against. The labor market has placed those five operations inside an AI-augmented workflow the pre-AI instruments were not built to see.

The instruments remain in place. The seven that §VI.1 named produce readings, the grades aggregate the readings, and the transcripts aggregate the grades. Every layer of the credentialing stack above the instrument continues to operate as if the reading at the bottom still correlates with the trait the labor market is hiring against. The correlation §VI.1 called construct-validity failure at the artifact layer runs up

the stack from there. Each aggregation inherits the reading underneath it. The reading at the bottom has shifted. The post-AI floor the labor market hires against is a floor the pre-AI instruments were never built to read.

VI.2.1 The William & Mary case at the instrument layer

William & Mary's Computing, Data Sciences, and Physics (CDSP) school is the nearest approximation in the dataset to an institutional answer to the measurement problem §VI.1 and §VI.2 named. The CDSP launch, announced in September 2024 with Douglas Schmidt as founding dean, was positioned explicitly around the "augmented intelligence" framing that pairs AI proficiency with liberal-arts critical thinking. The school's AI Power Hour series, co-hosted by the university CIO and the CDSP dean, is faculty-facing rather than integrity-facing, and the institutional press language around CDSP consistently treats AI as a curricular design question rather than a conduct question. The coding places W&M at Tier 3 — substantial readiness programs beside dominant governance — not because CDSP's curricular work is thin, but because it sits alongside a governance-dominant institutional posture that the curricular work has not displaced. The AI Policy Initiative produced its sixty-page report and recommended a permanent AI Innovation and Policy Council on the governance channel at the same time the CDSP work was forming on the curricular channel. The two tracks have not converged at the decision table where the institution sets its consolidated AI response.

The CDSP case reads against the measurement problem at two points. First, the curricular design work CDSP carries is the work the governance channel across the 53 institutions did not do. The instruments the governance channel owns levers to — the syllabus template, the integrity code, the detection contract — are not the instruments CDSP is building. CDSP's programming is upstream of the instruments at the course-design layer. Second, the case carries a caution. A former career center director at the institution who read the paper noted that the public record understates the depth of the institution's internal skepticism about whether the CDSP framing will travel past the school boundary into the fifty-plus other units where the curriculum actually runs. His reading, from inside the institution, is that the "augmented intelligence" formulation is a positioning claim the school's internal skeptics read as overreach. The CDSP case at the instrument layer is what institutional attention to the measurement problem looks like when it is present. What it does not resolve is whether the institution can act on the curricular reading at the scale the instrument layer requires across the other schools. The same public documentation that makes CDSP legible as a readiness investment does not document the institution's capacity to run the curricular reading across the eighty-odd programs outside CDSP where the instruments §VI.1 named will continue to produce readings against the pre-AI floor.

VI.3 The evidence from both sides of the transom

The §II evidentiary base extends past the three points §VI.2 lifted, and the mismatch holds across all of it. §VI.2 read three representative points from a single survey instrument and closed on the finding that the labor market has placed the five operations inside an AI-augmented workflow the pre-AI instruments were not built to see. Two additional survey tiers, a qualitative employer-interview record from one university's career office, and a body of graduate-side reporting all reach the same instruments at the bottom of the credentialing stack. At each tier the reading the instruments produce is one the labor market is no longer pricing.

The World Economic Forum's 2025 Future of Jobs report ranks AI and big data first among fastest-growing skill categories, with an 87 percent net increase over its five-year forecast horizon. Two-thirds of surveyed employers report plans to hire specifically for AI skills, and 39 percent of core workforce skills are projected as transformed or outdated by 2030. PwC's analysis of approximately one billion job advertisements found that skills in AI-exposed occupations are changing 66 percent faster than in other

roles, up from 25 percent the prior year. The instruments at the bottom of the credentialing stack measure capabilities the WEF respondents named as outdated and the PwC corpus showed turning over inside the same calendar window the instruments cycle students through. The reading produced by an essay or a problem set in 2026 is a reading of a capability whose half-life the labor market reports as compressing.

The Lehigh Center for Career and Professional Development's summer 2025 outreach to fifteen employer partners is the qualitative form the survey tiers do not capture. The five themes the CCPD team recorded described AI literacy as a baseline expectation rather than a differentiator and named human judgment, ethical use, and demonstrable application across non-technical fields as the competencies employers were screening for. One employer told the CCPD team that new hires are expected to be using AI tools by their third day on the job. Lehigh's provost summarized the finding as "applied fluency with AI, not model building." The third day is inside the orientation window. Applied fluency, as the provost named it, exists only in iteration with AI as a working partner, and the unit of production any of the seven instruments was built to measure is the student alone with the artifact.

The graduate-side evidence reports the same mismatch from the other side of the transom. Cengage's 2025 Graduate Employability Report recorded 30 percent of recent four-year graduates securing full-time work in their field of study, down from 41 percent the prior year. The eleven-point collapse occurred in twelve months, in the cohort the credentialing stack had just certified. Daniel Zhao, lead economist at Glassdoor, reframed the standard objection that graduate self-assessment is unreliable: the placement-rate collapse is the audit, and the self-reports are corroboration after the fact. The placement rate is a labor-market reading of what the credentialing stack produced, and the labor market is reading the same mismatch §VI.2 read at the instrument level. The reading at the bottom of the stack did not survive contact with the hiring decision.

The architects §V.5 named read the same labor market and reach the same instruments. Bowen and Watson's *Teaching with AI* names the AI skills salary premium, the 66 percent of business leaders who say they would not hire without AI skills, and the entry-level labor market collapse, and locates the entire institutional response in the classroom. The other AAC&U cluster sources reach the same routing from different distances: the Watson and Rainie leadership and faculty surveys exclude career services from the institutional-operations list, and the Anderson et al. student guide hands the workforce transition to the eighteen-year-old reader as a personal development project. The architect tier sees the labor market clearly and routes the response through the same governance levers and the same instruments §VI.1 named. The instruments the architect tier endorses by omission are the instruments §VI.1 and §VI.2 read as measuring pre-AI capabilities against a post-AI floor. The architects' reading becomes operational through those instruments.

The §II evidence and the §V.5 cluster are reading the same mismatch from two evidentiary positions. The labor market reports what it hires against. The architect tier endorses an institutional response that routes that signal through instruments the governance layer owns levers to. Both positions describe a single pattern, and the instrument is the point at which the mismatch becomes institutional fact. Every credentialing decision the stack makes is an aggregation of readings the instruments produced, and every aggregation inherits the reading underneath it. The credentialing stack is operating against a labor market both the §II evidence and the §V.5 cluster have already named. The mismatch becomes institutional fact at the instrument layer because the instrument layer is where the credentialing stack touches the student, and the student's degree records a reading the labor market has already stopped pricing at the floor the instrument was built against.

Section close The measurement problem is the form in which the governance response becomes institutional fact, and the instrument is the point of contact. The 9.7 percent §I named as the one figure

on the architects' declared-goals table they did not interpret is the share of the field's AI-engaged institutions that named the labor market as a goal at all. The other 90.3 percent named goals the governance layer owns the levers to act on, and every one of those goals routes through instruments built to read a pre-AI floor. §V.5 placed that number in architect-tier context and observed that the architects did not interpret it. §VI reads it at the instrument layer. Seeing the labor market clearly and routing the institutional response through instruments that cannot read it are the same act, and the act becomes operational the moment a student's degree records a reading the instruments produced against the floor §VI.2 read as gone. The 9.7 percent is the working definition of student education the field has endorsed by omission, and the instrument is the place that definition becomes the student's credential. §VII reads what the field built on top of the instrument layer and what it declined to build.

Section VII

THE OUTLIER

VII.1 Lehigh

VII.2 Independent-discipline convergence

VII. The Outlier

§VII reads what the field built on top of the instrument layer and what it declined to build. §VI named the measurement problem at the instrument layer where the governance response becomes institutional fact, and §VII reads the institutional fact at the case level. The 53-institution dataset contains one institution that built its AI program from the labor-market signal backward and fifty-two that did not. The asymmetry §III, §IV, §V, and §VI traced at the field level reaches its narrowest point at one institution.

§VII.1 walks Lehigh University as the case. Lehigh is the one institution of 53 that placed its career services unit at the decision table where the institution's AI strategy was set. The seat is the structural variable §VII.1 reads against the line-of-sight-without-standing finding §IV.2.4 named. §VII.2 reads the case alongside one independent commentary from a discipline that did not see the dataset and reached the same finding.

VII.1 Lehigh

Lehigh University placed its career services unit at the table where the institution's AI strategy was set, and the seat is the structural variable that turned line of sight into an institutional AI strategy. §IV.2.4 paragraph 5 named the variable at the constituency level; §VII.1 reads it at the case level. The position career services holds across the 53-institution dataset is the constituency closest to the labor-market signal, with the least standing in the channel where the institution decides what to do with the signal. That position came with a seat at Lehigh and not at the other fifty-two. The asymmetry §I framed as what the 9.7 percent built when it reached the institutional level reaches its narrowest case here: at the one institution, the seat existed.

The Center for Career and Professional Development is Lehigh's career services office. It sits outside the channels §III documented as the institution's standard AI response architecture and inside the channel where the institution sets AI strategy. CCPD originated the employer-outreach work, organized the findings, and brought them to the provost's office, which used them to frame Lehigh's institutional AI response. The institutional position the unit occupies is the variable. The fifty-two other career services units in the dataset have analytical capabilities equal to CCPD's, and the structural difference at Lehigh is procedural: CCPD has standing in the institution's AI decision-making channel.

Fifteen employers from a cross-section of industries sat for structured qualitative interviews conducted by CCPD staff over a three-week period in summer 2025.[95] The interviews produced five themed findings: AI use is broad and growing across industries; AI literacy is a baseline expectation rather than a bonus; human judgment and communication remain critical; ethical use of AI is a hiring signal; and AI readiness can be demonstrated by students outside of technical majors. The methodology is unremarkable in technique. The interviews ran on a five-question protocol covering entry-level skill expectations, differentiating competencies, emerging tools and use cases, power skills, and additional preparation gaps. Career services offices across the field have the institutional relationships, the interview infrastructure, and the analytical capacity to conduct an equivalent inquiry. What was distinctive at Lehigh was the routing. CCPD conducted the inquiry, organized the findings, and brought them upstream to the provost's office, which used the five themed findings as the specification from which

[95] See note 21.

the institutional AI response would be built. The interview record is, in §III's terms, the empirical input from which Lehigh's Tier 1 classification follows.

The Lehigh AI Fluency Framework took its curricular shape from the employer findings. Faculty AI development programming built out from the same source. The three-tier AI Readiness program structure followed. The College of Business designated the 2025–2026 academic year its "Year of AI@Work." Each component traces its origin to what employers described in the interview transcripts. The build direction at Lehigh is the inverse of the build direction §III documented across the other fifty-two institutions. The typical institutional AI response §III documented across the tiers is built from inside the institution outward, beginning with the governance committee's framing of the integrity question and proceeding through policy revision, faculty development programming, and whatever curricular guidance eventually reaches the student. The Lehigh response is built from outside the institution inward, beginning with the interview transcripts and proceeding through the framework, the AI Readiness program structure, the faculty programming, and the curricular content the student encounters. The reversal is the content of the structural-variable claim. The build direction follows from the seat.

Lehigh also opted out of Turnitin's AI detection product for the 2025–2026 academic year, citing false-positive rates and the concern that detection tools compromised student-teacher trust.[96] §III.3.5 documented the opt-out as the integrity-policy expression of the same institutional posture that produced the framework. §VII.1 presents it as evidence that the upstream structural variable shaped institutional decisions across the response stack: not only the curricular framework that originated from CCPD but the integrity policy that the governance committee at every other institution owned. The integrity layer at Lehigh was not insulated from the employer signal because career services at Lehigh was not insulated from the AI decision-making channel. Detection-tool adoption stood as the default integrity response under the structural conditions §III walked through across the other fifty-two institutions. The detection-tool opt-out is the case's secondary evidence; the framework is the headline.

Every institution in the 53-institution dataset has a career center. Every institution has employer relationships with the companies that hire its graduates. Lehigh is a private research university in Bethlehem, Pennsylvania, designated R1 by the Carnegie Classification in February 2025, with approximately 5,700 undergraduates and 7,500 total enrollment. Twenty-two of the other institutions in the dataset enroll more undergraduates; several enroll more than twice as many. The institutional infrastructure that produced the case is not unique to Lehigh's size, its type, or its resource base. The infrastructure exists at the other fifty-two institutions. What does not exist at the other fifty-two is the institutional structure that gives the career services unit standing in the channel where the institution makes AI decisions. The variable is procedural and institutional. It is the seat at the table.

VII.2 Independent-discipline convergence

Reza Barkhi and Michelle Seref published a Communication-format response in the *Journal of Macromarketing* in 2025, writing from Pamplin College of Business at Virginia Tech where both serve as associate deans.[97] The piece responds to Wided Batat's MECCDAL framework for marketing education

---

[96] See note 47.

[97] Reza Barkhi and Michelle Seref, "Holixec Education: Experiential and Personalized Learning Implemented with Mindfulness," Journal of Macromarketing 45(3), 2025, pp. 491–496. DOI: 10.1177/02761467251323513. Communication-format response to Wided Batat's MECCDAL/Holixec model articles (Batat 2024a; Batat 2024b). Both authors at Pamplin College of Business, Virginia Tech: Barkhi

and reaches the execution-readiness shift this paper reads from the 53-institution dataset. Barkhi and Seref do not engage the AI governance literature, do not name career services structure, and do not see the dataset. Their convergence is the section's evidence that the execution-gap finding is observable from disciplinary positions that share neither the paper's evidence base nor its analytical vocabulary.

Barkhi and Seref's argument runs from inside business school administration. They identify the disconnect between the traditional business curriculum, built before AI and automation became central, and the skills the present job market values. They locate the source of the misalignment in faculty incentive systems: department rankings tied to publications in elite disciplinary journals, departmental silos that resist interdisciplinary work, faculty evaluation built around structured coursework. Their prescription is that "business education must evolve to emphasize doing," a shift they frame as moving the human role toward execution as AI takes over data analysis. The diagnosis differs from this paper's at the field level. Barkhi and Seref locate the obstacle inside the faculty reward system; this paper locates it inside the governance channel that produced the institutional AI response. The two diagnoses do not overlap. What they share is the finding the diagnosis seeks to explain: that business education has not produced the execution readiness the labor market demands, and that the existing institutional response has not closed the gap. Two features of the convergence carry weight. The first is the disciplinary distance. Barkhi and Seref are responding to a marketing pedagogy paper from inside a decision and information science perspective; their argument is built without reference to the cross-institutional governance evidence this paper consolidates. The second is the institutional position. Both authors are associate deans at Pamplin, and Seref's portfolio explicitly includes career services and professional development. The pair co-chaired an AI taskforce at the college. They write from inside the policy-maker channel that §IV.2.4 documented as closed to career services across the 53-institution dataset, and from a portfolio that includes the constituency this paper identifies as the structural variable. The execution-gap finding lands in their commentary; the structural-variable claim does not. The convergence is on the finding, and the non-convergence on the diagnosis is a second piece of evidence: the structural-variable claim is hard to see even from the policy-maker position with career services in portfolio. §VII.1 reads the case where the structural variable was named and acted on; §VII.2 reads an independent commentary where the same finding lands without the variable being named.

§VII reads the asymmetry §I framed at its narrowest case. §III walked the asymmetry through the tiers at the field level, §IV walked it through the constituencies' working life, §V placed it against the literatures, and §VI named the measurement problem at the instrument layer where the governance response becomes institutional fact. At the case level, the asymmetry has a structural cause. The variable is procedural and institutional. §VII.1 named the variable at the case where it was acted on, and §VII.2 named the variable's force where it was invisible. The case at §VII.1 names the variable in the position where it produced the outcome, a build direction that started from the labor-market signal and worked back into curriculum, faculty programming, and the integrity-policy posture; the convergence at §VII.2

---

as Associate Dean of Strategy, Global Initiatives, and Innovations and KPMG Professor of Accounting and Information Systems (per author biography in the source); Seref as Associate Dean of Undergraduate Programs and Collegiate Professor of Business Information Technology Department, with portfolio coordinating undergraduate business curriculum, academic advising, undergraduate teaching assistantships, undergraduate research, career services and professional development, undergraduate international programs, living learning communities, and student organizations and mentorship programs (per author biography in the source). Both authors co-chaired an AI taskforce at Pamplin per the author biographies. Direct quote in §VII.2 ("business education must evolve to emphasize doing") drawn from the article abstract; tighter seven-word lift from the longer abstract sentence "We agree that business education must evolve to emphasize doing—encouraging students to apply an execution mindset and leverage AI to analyze data."

names the same outcome in a position where the variable did not surface, an execution-gap diagnosis that lands at faculty incentive structures with the structural-variable claim never named. The two readings are non-overlapping evidence of the same finding from different vantages, not parallel walks of the same evidence. The asymmetry §I framed reaches its narrowest case at one institution that had the seat and at fifty-two that did not.

Section VIII

THE DISPOSITION

VIII. The Disposition

The finding names an asymmetry the field has not yet answered for. Governance infrastructure at the institutional level exists to address an assessment-and-policy problem the field has defined; readiness infrastructure at the graduate level does not exist because the field has not defined the question the labor market is asking. The career services unit is the variable §VII.1 made load-bearing: the position that sits at the edge of the employer signal and whose presence at the AI decision table was the structural difference between the one institution that built its AI program from the labor-market signal backward and the fifty-two that did not. Governance bodies at the fifty-two receive the employer signal through the channels that exist at the institution, but those channels do not reach the AI decision table. Accreditation categories were fixed before the question the labor market is asking existed. The career services unit sits outside the AI decision channel because the institutional structure routes AI decisions through other positions. The finding puts the structural question in front of accreditors, boards of visitors, governance committees, and the career services units themselves: whether the channels through which each body sees the institution include the labor-market signal, and whether the decision table at which each body has standing to act includes the position from which the signal enters. None of those bodies has taken up the question at scale. The asymmetry is the residue of institutional decisions each body is positioned to revisit, not an interpretive problem further analysis will resolve.

The pattern this paper documented is the outcome of institutional decisions the governance layer was equipped to make and the readiness question the governance layer was not equipped to answer. The ratio is the paper's specification of that outcome at the field's current scale, documented across fifty-three institutions through the spring of 2026. The governance-to-readiness gap has been the architecture of the field since 2023, not an artifact of the current moment. The institutional decisions that produced the ratio are still being made. The asymmetry sits inside the channels the governance layer owns, not outside institutional control. What the paper cannot resolve is whether the Lehigh case will remain exceptional or become the baseline; what the four bodies can act on is the institutional decisions that make the answer to that question turn on their standing. The ratio is an outcome the field chose. The variable is a variable the field can reposition.

Institutional Proximity and Scope Disclosures

Institutional proximity. The author is an independent researcher based in Williamsburg, Virginia, and has no employment, consulting, or governance relationship with William & Mary. The 53-institution profiling framework was applied uniformly. William & Mary receives more detailed treatment in the institutional analysis than most institutions in the dataset. This reflects two factors: the Virginia Reference Guide, which profiles all 30 Virginia institutions at greater depth than the national peer benchmark, and the public availability of W&M governance documentation — particularly the AIPI process, the CDSP launch, and Board of Visitors materials — which is more extensive than what most institutions publish. The deeper profile is a function of source availability, not evaluative emphasis. William & Mary is classified in this analysis as Tier 3 (readiness programs present, governance dominates), not as an outlier.

Limitations and Future Research

The dataset is the most detailed independent institutional analysis of AI governance in American higher education, profiling 53 institutions across ten governance and readiness dimensions using public-source documentation. Broader surveys exist — EDUCAUSE's 2024 and 2025 AI Landscape Studies surveyed over 900 technology professionals across hundreds of institutions — but they capture self-reported policy status through individual respondents rather than producing institution-level profiles with documented evidence. The distinction matters: this study's value is not breadth of sample but depth and consistency of institutional documentation. It still carries constraints worth naming plainly.

Public-source methodology. All institutional data was collected from publicly accessible materials: institutional websites, policy documents, press releases, official announcements, academic catalogs, student and faculty handbooks, library research guides, IT services pages, teaching center archives, and campus news coverage. No internal documents, interviews, non-public data, or behind-authentication materials were used. The data captures what institutions have chosen to publish, not the full scope of what they are doing. This constraint cuts both ways: workforce readiness activity may be more extensive than the data shows, occurring through internal channels and individual faculty efforts that do not reach public websites. Governance activity may be more performative than the data shows, with published policies that exist on paper but do not meaningfully shape institutional practice. The loose coupling framework suggests both possibilities are likely. The analysis treats absence of public documentation as an open question, not a conclusion.

Evidence window. The institutional profiles reflect publicly documented activity as of March 2026. The generative AI landscape is evolving rapidly. The employer demand data spans January 2023 through March 2026, with concentration in 2024 and 2025. The specific data points in this paper have a limited shelf life. The structural patterns — the governance-first default, the readiness gap, the theoretical explanations for why these patterns emerge — are likely to persist longer than any individual policy or survey finding.

Employer data limitations. The demand-side evidence carries known limitations. The Microsoft Work Trend Index surveys knowledge workers, excluding large segments of the labor force. McKinsey's surveys reflect senior executive views at large organizations. NACE's samples are small and skew toward large companies. The Cengage and CarringtonCrisp surveys are produced by organizations with commercial interests. The Lehigh outreach, while unique in the dataset, is a qualitative thematic summary of 15 employer conversations, not a quantitative instrument. There is a documented discrepancy between employer sentiment (66 percent of leaders say they will not hire without AI skills) and job posting behavior (10.5 percent of entry-level descriptions explicitly mention AI). No single survey carries the argument. The convergence across independent sources is what makes the demand-side evidence persuasive, but the magnitude of the signal may be overstated by any individual instrument.

Sample composition. The 53 institutions are not a random sample of American higher education. The national peer group is weighted toward selective, well-resourced institutions. The Virginia dataset provides institutional type diversity within one state but not geographic diversity. Community colleges, which enroll nearly half of all American undergraduates, are not included. Joseph Fuller of Harvard Business School warned at a November 2025 conference that without creating an AI-literate generation across the income distribution, a new digital divide will emerge. The patterns documented here may not generalize to all institutional types, regions, or sectors.

Causation and outcomes. This paper documents a pattern. It does not demonstrate that governance-heavy institutions produce less-prepared graduates or that readiness-oriented institutions produce better-prepared ones. That causal question requires outcome data — graduate employment rates, employer

satisfaction measures, alumni self-assessments of AI competency, longitudinal career tracking — that this study does not collect.

Future research directions. Five lines of inquiry follow directly.

First, outcome measurement: linking institutional AI governance posture to measurable workforce outcomes.

Second, longitudinal tracking: repeating the institutional analysis at twelve-month intervals to determine whether the governance-first pattern is stable or evolving.

Third, the "instructor decides" model: measuring how many courses in a typical undergraduate program permit AI use, how many prohibit it, and how much variation exists within a single institution.

Fourth, institutional type analysis: testing whether the governance-readiness relationship varies across R1 research universities, selective privates, regional publics, small liberal arts colleges, and community colleges.

Fifth, international comparison: testing whether the governance-first default holds in different regulatory environments and labor market structures.

This paper documents how American higher education responded to generative AI in the first three years after its widespread public availability. The response was structurally predictable, theoretically explicable, and measurably misaligned with the labor market these institutions serve. The graduates walking out the door this spring will discover the distance between what their institutions governed and what their employers expect. Whether that distance closes — and how fast — depends on decisions that most universities have not yet made.

About the Author

Tommy N. Turner is an independent researcher based in Williamsburg, Virginia. He holds an MBA from Emory University's Goizueta Business School and a certificate from the Harvard Bok Center for Teaching and Learning. Over thirty years of senior leadership experience in banking, technology, and consulting inform this research. His work examines how institutions respond to technological disruption, with particular attention to the gap between organizational governance and workforce outcomes. This paper draws on original institutional profiling of 53 American universities conducted between October 2025 and March 2026.

Correspondence: tommy@gofast.pro

Suggested Citation

Turner, Tommy N. "Governance Without Readiness: How 53 American Universities Responded to Generative AI." Working Paper, Version 6.3. Williamsburg, VA: October 2026. Zenodo. https://doi.org/10.5281/zenodo.23442320

Companion Materials

This working paper is accompanied by one companion dataset, deposited separately:

53-Institution Coding Table (Excel). Zenodo, https://doi.org/10.5281/zenodo.21284051. Institution-level coding across four governance variables and five readiness variables for all institutions in the dataset, with summary counts, tier classification, a source register and the codebook.

Appendix A: 53-Institution Coding Table

Institution	Sample	Type	Cmte	Syll.	Integ.	Detect.	Data Cl.	AI Deg.	Empl.	Career	Enterp.	Fac. Dev.	Gov.	Ready	Tier
Lehigh University	Peer	Private	Y	N	Y	Out	Y	Y	Y	Y	Y	Y	High	High	1
Clemson University	Peer	Public	Y	Y	Y	Caut.	Y	Y	Y	Y	Y	Y	High	High	2
Rice University	Peer	Private	Y	Y	Y	Caut.	Y	Y	N	N	Y	Y	High	High	2
Vanderbilt University	Peer	Private	Y	Y	Y	Out	Y	Y	N	N	Y	Y	High	High	2
University of Georgia	Peer	Public	Y	Y	Y	On	Y	Y	N	N	Y	Y	High	High	2
Marymount University	Virginia	Private	Y	N	Y	Ment.	N	Y	N	N	Y	Y	Medium	High	2
Northeastern University	Peer	Private	Y	N	Y	Caut.	Y	Y	N	N	Y	Y	High	High	2
George Mason University	Virginia	Public	Y	Y	Y	Caut.	Y	Y	Y	N	Y	Y	High	High	2
William & Mary	Virginia	Public	Y	Y	Y	Caut.	Y	Y	N	N	Y	Y	High	Medium	3
Brown University	Peer	Private	Y	N	Y	On	Y	Y	N	N	Y	Y	High	Medium	3
Duke University	Peer	Private	Y	N	Y	On	Y	Y	N	N	Y	Y	High	Medium	3
University of Southern California	Peer	Private	Y	N	Y	Caut.	Y	Y	N	N	Y	Y	High	Medium	3
Georgetown University	Peer	Private	Y	N	Y	Caut.	Y	Y	N	N	Y	Y	High	Medium	3
American University	Peer	Private	Y	Y	Y	On	Y	Y	N	N	Y	Y	High	Medium	3
Emory University	Peer	Private	Y	N	Y	Caut.	Y	Y	N	N	Y	Y	High	Medium	3
Stony Brook University	Peer	Public	Y	N	Y	On	Y	Y	N	N	Y	Y	High	Medium	3
Virginia Tech	Virginia	Public	Y	Y	Y	Out	Y	N	N	N	Y	Y	High	Medium	3

Institution	Sample	Type	Cmte	Syll.	Integ.	Detect.	Data Cl.	AI Deg.	Empl.	Career	Enterp.	Fac. Dev.	Gov.	Ready	Tier
University of Virginia	Virginia	Public	Y	Y	Y	Caut.	Y	N	N	N	Y	Y	High	Medium	3
Virginia Commonwealth Univ.	Virginia	Public	Y	N	Y	Caut.	Y	Y	N	N	Y	Y	High	Medium	3
Washington & Lee University	Virginia	Private	Y	N	Y	Caut.	N	Y	Y	N	Y	Y	Medium	Medium	3
Dartmouth College	Peer	Private	Y	N	Y	On	Y	N	N	N	Y	Y	High	Low	4
Boston College	Peer	Private	Y	N	Y	Caut.	Y	N	N	N	Y	Y	High	Low	4
Brandeis University	Peer	Private	Y	N	Y	Caut.	Y	N	N	N	Y	Y	High	Low	4
George Washington University	Peer	Private	Y	N	Y	Caut.	Y	N	N	N	Y	Y	High	Low	4
Santa Clara University	Peer	Private	Y	Y	Y	On	Y	N	N	N	Y	Y	High	Low	4
Tufts University	Peer	Private	Y	N	Y	Caut.	Y	N	N	Y	Y	Y	High	Low	4
Univ. of North Carolina Chapel Hill	Peer	Public	Y	Y	Y	On	Y	N	N	N	Y	Y	High	Low	4
Villanova University	Peer	Private	Y	N	Y	Caut.	Y	N	N	N	Y	Y	High	Low	4
Wake Forest University	Peer	Private	Y	Y	Y	Caut.	Y	N	N	N	Y	Y	High	Low	4
Washington Univ. in St. Louis	Peer	Private	Y	N	Y	Caut.	Y	N	N	N	Y	Y	High	Low	4
Old Dominion University	Virginia	Public	Y	Y	Y	Caut.	Y	Y	N	N	N	Y	High	Low	4
James Madison University	Virginia	Public	Y	Y	Y	Caut.	Y	N	N	N	Y	Y	High	Low	4
Christopher Newport University	Virginia	Public	Y	N	Y	Caut.	Y	N	N	N	Y	Y	High	Low	4
University of Mary Washington	Virginia	Public	Y	Y	Y	Caut.	Y	N	N	N	Y	Y	High	Low	4
Liberty University	Virginia	Private	Y	N	Y	Caut.	Y	N	N	N	Y	Y	High	Low	4
Radford University	Virginia	Public	Y	N	Y	On	Y	Y	N	N	N	Y	High	Low	4
University of Richmond	Virginia	Private	Y	N	Y	Caut.	N	N	N	N	Y	Y	Medium	Low	4

Institution	Sample	Type	Cmte	Syll.	Integ.	Detect.	Data Cl.	AI Deg.	Empl.	Career	Enterp.	Fac. Dev.	Gov.	Ready	Tier
Hampton University	Virginia	Private	Y	N	Y	Caut.	N	Y	N	N	Y	N	Medium	Low	4
Shenandoah University	Virginia	Private	Y	Y	Y	Caut.	N	Y	N	N	Y	Y	Medium	Low	4
Norfolk State University	Virginia	Public	Y	N	Y	On	Y	N	N	N	Y	N	Medium	Low	4
Longwood University	Virginia	Public	Y	Y	Y	Caut.	N	N	N	N	Y	Y	Medium	Low	4
Virginia Military Institute	Virginia	Public	Y	N	Y	On	Y	N	N	N	N	N	Medium	Low	4
Regent University	Virginia	Private	Y	Y	Y	On	N	N	N	N	N	N	Medium	Low	4
Virginia State University	Virginia	Public	Y	N	Y	On	N	Y	N	N	N	N	Medium	Low	4
University of Lynchburg	Virginia	Private	Y	N	Y	Caut.	Y	Y	N	N	Y	N	Medium	Low	4
Roanoke College	Virginia	Private	Y	N	Y	On	Y	N	N	N	N	N	Medium	Low	4
Hampden-Sydney College	Virginia	Private	Y	Y	Y	Caut.	Y	N	N	N	N	Y	Medium	Low	4
Virginia Wesleyan University	Virginia	Private	Y	Y	Y	Caut.	Y	N	N	N	Y	N	Medium	Low	4
Randolph-Macon College	Virginia	Private	Y	N	Y	Caut.	Y	N	N	N	N	N	Medium	Low	4
Emory & Henry University	Virginia	Private	N	N	N	N/D	N	N	N	N	Y	N	Low	Low	5
Hollins University	Virginia	Private	Y	N	N	Caut.	N	N	N	N	Y	N	Low	Low	5
Sweet Briar College	Virginia	Private	Y	N	N	Caut.	N	Y	N	N	Y	N	Low	Low	5
UVA College at Wise	Virginia	Public	Y	N	N	Caut.	N	N	N	N	Y	N	Low	Low	5

Y = documented in public sources. N = not documented. Detection: Out = opted out or rejected; Caut. = available with caution; On = enabled; Ment. = mentioned but no clear policy; N/D = not documented. Gov./Ready Level: High/Medium/Low, a judgment informed by the number of documented indicators. Tier: whole-profile classification (§III.3); 1 = employer-originated; 2 = substantial readiness + governance; 3 = readiness present, governance dominates; 4 = governance-dominant; 5 = minimal documented activity. Variable definitions and coding rules in the companion dataset.

Appendix B: Voices by Constituency

Selected direct statements from named individuals cited in this paper, organized by constituency. Each quote is drawn from a publicly documented source cited in the endnotes.

Employers and Hiring Managers

Name / Title	Statement	Source
Brendan Humphreys *CTO, Canva*	"We evaluate the interactions with the AI as much as the output."	Business Insider, February 2026
Alejandro Castellano *CEO, Caddi*	"We want to see how they work in real conditions."	Fortune, June 2025
Alex Salazar *Cofounder, Arcade*	"If workers will use AI on the job, it makes no sense to create an artificial test that does not reflect real work."	Business Insider, February 2026
Ehsan Mokhtari *CTO, ChargeLab*	"Candidates were avoiding AI tools during interviews because they assumed they would be penalized."	Fortune, June 2025
Omar Haroun *CEO, Eudia*	"A recent graduate with zero experience, obsessed with AI, achieved 10 times the output of a more experienced employee working without the boost from the new technology."	Inc., November 2025
Nolan Church *Former Head of Talent, Google and DoorDash*	"Companies used to help us mature. Now they say, 'No, we need you to add value on Day One.' Onboarding and training? We don't even think about those things anymore."	Riviera Partners podcast, March 2026
Alan Price *Global Head of Talent Acquisition, Deel*	"Most interviews don't get past level one assessment of candidates, or not test at all."	Deel / LinkedIn, March 2026
Alan Price *Global Head of Talent Acquisition, Deel*	"A lot of organizations have stopped [case studies] because people use AI, and you can't tell if you're testing their expertise or their prompting ability."	Deel, March 2026
Daniel Zhao *Lead Economist, Glassdoor*	"AI doesn't replace the professional's judgment, but universities are often so focused on the 'risk' of AI that they aren't teaching the 'audit' of it."	Fortune, 2026
Jonna Sercombe *CEO, Cegos*	"If a candidate tells me they haven't used AI because their university 'didn't allow it,' they aren't showing integrity. They're showing a lack of preparation for the modern workforce."	LinkedIn, 2026
Nico McCord *Recruiter, Ten10*	"It felt like I met a robot, not a person."	LinkedIn, March 2026
Nathalie Agnew *Managing Director, Muckle Media*	"We received 50-plus applications in two days and rejected most because they were transparently produced by ChatGPT with no personality."	LinkedIn, March 2026

Faculty

Name / Title	Statement	Source
Elanah Uretsky *Associate Professor, International and Global Studies, Brandeis University*	"I'm not convinced AI teaches critical reasoning. I don't use AI in my classes at all, and I don't let my students use it on assignments. If we can't think, we won't succeed as a society."	Brandeis Magazine, Winter 2024/2025
Aldo Musacchio *Professor of Business, Brandeis International Business School*	"I've used AI very successfully in our entrepreneurship classes. AI tools can help students go through the prototyping process faster than ever before."	Brandeis Magazine, Winter 2024/2025
Marc Watkins *Academic Innovation Fellow, University of Mississippi*	"The same professors forecasting academic doom were now giddily rebranding themselves as 'AI-ready educators.' The initial panic about plagiarism gave way to a resigned embrace."	Current Affairs, December 2025
Sarah Rose Cavanagh *Senior Associate Director for Teaching and Learning, MIT*	"Instructors must balance compassionate challenge with the reality of AI, warning against passive leniency while acknowledging that student overreliance can worsen student learning outcomes."	MIT Teaching + Learning Lab, 2025
Anna Mills *English Faculty, Canada College*	"Institutions cannot allow the value of those degrees to disappear."	OneHE, January 2026
Lin Zhong *Professor of Computer Science, Yale University*	"The curriculum is stuck in the 1980s. If your students don't learn how to use AI, they lose their job."	The New Journal at Yale, 2025
Frederique Laubepin *Faculty, University of Michigan School of Public Health*	"I knew I wanted to utilize AI in a way that harnessed its ability to solve actual problems, not just serve as an academic exercise."	Michigan Public Health, May 2025
Helen Choi *Associate Professor of Technical Communication, USC*	"My job is not to count essays."	Daily Trojan, April 2026
Leslie Clement *Professor, Johnson C. Smith University*	"We encourage students to use it because we know they're going to use it, but to use it in a responsible way. We want them not only to use the tools for good but also to interrogate them."	WFYI/NPR, March 2026
Bethany Nowviskie *Dean of Libraries, James Madison University*	"What is the impact of AI on our ability to foster our students' critical thinking and everybody's collective production of new knowledge?"	WHSV, July 2024
Gene Roche *Executive Professor Emeritus, William & Mary School of Education; special advisor to the dean of the Mason School of Business on AI; former career center director*	"Professional staff who have deep knowledge but are not afforded the prestige of faculty."	Personal communication, April 9, 2026

Students

Name / Title	Statement	Source
Hannah Elder *Student, University of North Carolina*	"Faculty should incorporate AI into curricula so students learn how to use it responsibly, instead of leaving them to figure it out alone or treating it only as a cheat code."	WFYI/NPR, March 2026
Kelsey Auman *Student, University at Buffalo*	"Unfortunately, with Turnitin, the numbers don't mean anything. There's no reasoning given. We need the university to realize things have been done wrong."	Spectrum News, May 2025
Haishan Yang *Ph.D. Student, University of Minnesota Twin Cities*	"A death penalty."	MPR News, January 2025
"Avery" *Freshman, Louisiana State University*	"They're using AI to detect AI, which makes no sense."	WBRZ, January 2026
"Edward" *Student, Yale University*	"Everyone was freaking out."	The New Journal at Yale, 2025
Zhariyah Laughlin *Student, University of Michigan School of Public Health*	"Working on the app and seeing firsthand how AI can transform public health interventions completely shifted my perspective. AI is an essential part of the work I want to do."	Michigan Public Health, May 2025
Boston College Student AI Advisors *Student Advisory Group, Boston College*	"Students are eager to move past the dominant institutional narrative that suggests that using AI is primarily an issue of academic integrity. They want more nuanced ethical guidance, practical skill development, and open dialogue."	Boston College CDIL Report, 2025

Career Center Professionals and Institutional Leaders

Name / Title	Statement	Source
May Wilson *Assistant Director, UC Berkeley Career Engagement*	"AI is the big elephant in the room. Employers will probably expect students to have practical AI-use skills, such as prompting AI to review material or identify domain trends."	UC Berkeley Graduate Division, February 2026
Bob Beaudet *Director of Employer Relations, St. John's University*	"Employer partners are conveying their expectations clearly that graduates should possess AI fluency, specifically in communications and business analytics roles."	St. John's University, 2025
Irma Becerra *President, Marymount University*	"The speed of AI adoption is unprecedented, and the question for higher education is, 'can we keep up?'"	Inside Higher Ed U.S. Universities Summit, 2025
Joseph Helble *President, Lehigh University*	"AI should free more time for the deeply human parts of work."	Lehigh University Town Hall, November 2025
Douglas C. Schmidt *Dean, CDSP, William & Mary*	"Augmented intelligence — the combination of AI proficiency with liberal arts critical thinking."	William & Mary Tack Faculty Lecture, Spring 2025
Joseph Fuller *Harvard Business School Project on Workforce*	"Without creating an AI-literate generation across the income distribution, a new digital divide will emerge."	Harvard Gazette, November 2025
Reza Barkhi and Michelle Seref *Associate Deans, Pamplin College of Business, Virginia Tech*	"Business education must evolve to emphasize doing."	*Journal of Macromarketing* 45(3), 2025