Mapping the Structural Divide: Institutional Resilience, Post-College Market Position, and Artificial Intelligence Exposure Across U.S. Higher Education

Kyle Saunders

Department of Political Science, Colorado State University

Author Note

Kyle Saunders, Department of Political Science, Colorado State University, Fort Collins, CO 80523.

Correspondence: kyle.saunders@colostate.edu

Abstract

I present a framework for mapping 1,556 U.S. four-year colleges and universities along two composite dimensions, Institutional Resilience and Post-College Market Position, using publicly available federal data. Post-College Market Position captures whether graduates enter career pathways that are economically viable and durable. The analysis documents three findings. First, American higher education exhibits sharp tier-level stratification: 82% of R1 institutions fall in the High Capacity quadrant, while 50% of baccalaureate and 43–46% of master's institutions fall in High Stress. Second, 37% of institutions are simultaneously shrinking in enrollment and scoring below the median on resilience, concentrated in baccalaureate and smaller master's programs, a bifurcation pattern consistent with Grawe's (2018) demand divergence thesis. Third, as an exploratory extension, a novel institution-level measure of AI-related labor market exposure shows that theoretical AI task exposure and observed real-world AI adoption are approximately uncorrelated at the field level ($\rho \approx -0.09$), suggesting two distinct timelines of potential disruption. The degree fields most exposed to AI disruption are currently the highest-earning fields for new graduates ($\rho = 0.257$), meaning disruption at scale would affect the career pathways with the highest current economic returns. The complete dataset, replication materials, and an interactive tool are publicly available.

Keywords: higher education, institutional resilience, AI exposure, labor market outcomes, institutional stratification, demographic cliff, strategic positioning

1. Introduction

American higher education is stratified in ways that are well documented but incompletely measured. The institutional hierarchy that structures the sector, from well-resourced research universities with deep endowments and global brand recognition to tuition-dependent regional colleges serving local labor markets, has been a central concern of higher education scholarship for decades (Bastedo, 2012; Morphew & Huisman, 2002). What has changed is the convergent nature of the pressures now acting on this hierarchy. Enrollment has declined at most institution types since 2010, with total undergraduate enrollment falling 15% from 18 million to 15.4 million by 2022 (NCES, 2023), and the demographic contraction predicted by WICHE's projections is no longer a forecast: high school graduate numbers will peak in 2025 and decline an estimated 13% by 2041, with the steepest losses in the Northeast and Midwest (Bransberger et al., 2024). On the revenue side, tuition discount rates at private nonprofits exceed 50% (NACUBO, 2023), tuition growth has been at or below inflation since 2018 (Ma & Pender, 2024), and pandemic-era relief funds that temporarily masked structural deficits have been fully spent (Kelchen et al., 2025).

These pressures are unevenly distributed, and that unevenness is the central empirical fact this paper documents. The question is not whether American higher education is under stress but how that stress distributes across the institutional landscape, what structural characteristics predict positioning, and whether emerging pressures from artificial intelligence are introducing dimensions of risk that existing diagnostic tools do not capture.

Literature Review and Conceptual Framing

Four bodies of scholarship inform this analysis.

Institutional stratification and organizational fields. The structure of American higher education has long been characterized by steep institutional stratification, in which a small number of resource-rich institutions occupy positions of enduring advantage while the majority operate under materially different conditions (Bastedo, 2012). Slaughter and Rhoades (2004) documented how market-oriented behaviors (program expansion into professional fields, revenue diversification through auxiliary enterprises, competitive positioning for tuition-paying students) have intensified across the sector, blurring traditional mission boundaries while reinforcing resource hierarchies. Morphew and Huisman (2002) showed that institutional diversity, often invoked as a distinguishing virtue of American higher education, coexists with powerful isomorphic pressures that push institutions toward similar programmatic offerings regardless of their structural position. Jaquette and Curs (2015) demonstrated that enrollment management strategies, particularly aggressive out-of-state recruitment by public universities facing state appropriation declines, have created competitive dynamics that cascade through enrollment markets, affecting the regional institutions least equipped to compete on brand or amenities. The framework I develop here offers a way to locate institutions within this stratified landscape using outcome-based measures rather than reputational proxies, and to identify where compound pressures are most concentrated.

Demographic and enrollment pressures. Grawe (2018) provided the foundational analysis linking declining birth rates to future enrollment pressure, and his Higher Education Demand Index demonstrated that elite institutions would face rising demand even as regional and less selective institutions experienced substantial contraction, a bifurcation pattern the present analysis documents empirically. Grawe's (2021) follow-up confirmed this divergence and introduced the concept of institutional agility as a moderating factor. The WICHE projections

(Bransberger et al., 2024), now in their 11th edition, provide state-level granularity that I incorporate as a framework component. Hillman (2014, 2016) showed that state funding decisions compound demographic pressures: institutions in states with declining populations and tightening budgets face simultaneous enrollment and revenue contractions, creating a structural trap that institutional strategy alone cannot escape. Tandberg (2010, 2013) demonstrated that state higher education funding is shaped by political dynamics (gubernatorial priorities, legislative composition, interest group competition) that operate independently of institutional need. My contribution is the integration of demographic projections into a composite assessment that also accounts for financial resilience, labor market outcomes, and AI exposure, offering a multidimensional picture of institutional positioning that no single indicator captures.

Institutional financial health and closure prediction. Kelchen, Ritter, and Webber (2025) assembled the most comprehensive dataset to date on 8,633 U.S. postsecondary institutions and developed models that predict closures with 83% accuracy, substantially exceeding the 77% accuracy of federal financial responsibility composite scores. Zemsky, Shaman, and Baldridge (2020) developed the widely-adopted "College Stress Test" emphasizing enrollment market position and net tuition revenue as leading indicators of financial distress. Galloway (2020) positioned approximately 436 institutions on a 2×2 of institutional value versus vulnerability, demonstrating that the sector could be mapped as a strategic positioning exercise; his methodology relied on brand-perception proxies and a COVID-specific vulnerability measure. The present analysis extends this literature by adding a post-college market position dimension that captures a different kind of vulnerability: an institution can be financially solvent but positioned in a labor market that is contracting or being restructured, creating risks that manifest on a longer timescale than financial distress.

AI and labor market disruption. Eloundou et al. (2023) estimated the share of occupational tasks exposed to large language models. Felten, Raj, and Seamans (2023) measured occupational heterogeneity in AI exposure. Most recently, Massenkoff and McCrory (2026) introduced "observed exposure," combining theoretical AI capability with actual usage data from the Anthropic Economic Index, and documented that actual adoption is a fraction of theoretical capability. They found suggestive evidence that hiring of younger workers (ages 22–25) has slowed in exposed occupations, consistent with entry-level displacement preceding broader effects. Karger et al. (2026) surveyed economists and AI researchers on AI's economic effects, finding that under rapid progress scenarios, labor force participation could decline to 55% by 2050 (approximately 10 million fewer jobs attributable to AI), with the wealthiest 10% forecast to hold 80% of total wealth. I extend this literature by translating occupation-level AI exposure into higher education institution-level vulnerability through a pipeline that weights for entry-level career pathways, connecting labor market disruption research to the institutional stratification literature in a way that neither body of work has done independently.

Contribution

This paper makes three contributions and one exploratory extension. First, I construct a transparent, reproducible framework for institutional positioning using exclusively public data, offering an alternative to both proprietary consulting analyses and the Galloway framework's reliance on brand-perception proxies. The framework is designed for institutional leaders, enrollment strategists, and policy researchers, but also to facilitate broader conversation about the structural challenges the sector faces. Second, I document pronounced tier-level stratification and a striking enrollment × resilience interaction: nearly 40% of institutions are simultaneously shrinking and below-median on resilience, concentrated in specific Carnegie tiers and geographic regions. Third, I provide a complete public dataset, replication infrastructure, and an interactive tool (available at https://kylesaunders.com/university-map) enabling institutional-level exploration; the tool has received more than 200,000 page views since its launch.

As an exploratory extension, I develop what is, to my knowledge, the first higher education institution-level measure of AI-related labor market exposure, building from O*NET task-level analysis through an entry-level weighting and CIP-SOC crosswalk pipeline. When benchmarked against observed AI adoption data from the Anthropic Economic Index, this measure shows that theoretical task exposure and actual adoption are nearly uncorrelated at the field level, consistent with the possibility of two distinct disruption timelines. I present this component as a provisional contribution intended to open a line of inquiry rather than a fully validated measure.

2. Data and Institutional Universe

The analysis draws on seven public data sources: IPEDS institutional characteristics, finance, enrollment, and completions data (2024); College Scorecard institution-level outcomes (most recent cohort); WICHE high school graduate projections (11th edition); the O*NET occupational database (v29.0); the NCES CIP-SOC crosswalk (2020); the Anthropic Economic Index (August 2025); and Census PSEO earnings data. Complete variable definitions are in Supplementary Appendix A.

Institutions were included if they were four-year, degree-granting, public or private nonprofit, and currently operating. For-profit institutions were excluded because they operate under a fundamentally different regulatory framework, face distinct accountability requirements, and report to IPEDS using conventions that limit comparability on several key variables. This yielded 1,609 institutions (600 public, 1,009 private nonprofit). Of these, 1,556 have sufficient data to compute both composite scores and receive quadrant assignments. Institutions are classified using the 2025 Carnegie Classification system as reported in IPEDS. I derive a 9-tier grouping variable using the Research Activity Designation as the primary axis for research institutions (R1, R2, RCU) and Award Level Focus for non-research institutions.

A critical data-processing correction was required for the IPEDS finance files. Public institutions report under GASB standards while private nonprofits report under FASB standards. The GASB form uses cumulative revenue fields that, if misinterpreted, produce tuition dependence ratios near 1.0 for all public institutions. The corrected approach uses the appropriate disaggregated fields; after correction, mean tuition dependence is 16.7% for publics and 77.8% for private nonprofits, in line with established patterns in higher education finance (details in Supplementary Appendix B).

Of the 1,609 institutions meeting inclusion criteria, 53 (3.3%) lack sufficient data for both composite scores. Missingness is not random: it is concentrated among smaller private institutions, Baccalaureate and Special Focus tiers, and institutions that do not report to the Scorecard on all outcome variables. The mapped universe therefore somewhat underrepresents the smallest and most resource-constrained institutions — the population most likely to be structurally vulnerable.

3. Framework Construction

X-Axis: Institutional Resilience

Institutional Resilience is a composite of four components, each converted to a percentile rank (0–1) within the sample before aggregation with equal weights:

Endowment per student. End-of-year endowment market value divided by FTE undergraduate enrollment. Captures institutional financial buffer. Structurally advantages wealthy private institutions, partially offset by the revenue diversification component.

Revenue diversification. One minus the ratio of net tuition revenue to total core revenues. Measures the degree to which an institution depends on a single revenue stream. Advantages public institutions with state appropriations and research funding.

Enrollment trajectory. Percent change in 12-month unduplicated headcount from 2019 to 2024. Captures five-year enrollment momentum spanning the pandemic disruption and subsequent recovery.

Selectivity. One minus the admission rate. A demand signal and brand proxy. The most contested component, as it disadvantages broad-access institutions serving important equity missions. Results under a "no selectivity" specification are reported in the sensitivity analysis.

Y-Axis: Post-College Market Position

The second axis captures whether an institution's graduates enter career pathways that are economically viable and durable, and therefore whether the institution's value proposition to prospective students is structurally sound. This matters for institutional strategy because labor market outcomes feed back into enrollment demand: prospective students and their families increasingly evaluate institutions on post-graduation employment and earnings (Hillman, 2014), and institutions whose graduates enter weak or restructuring labor markets face downstream risks to enrollment, alumni giving, and public support for continued investment. The axis is a hybrid: it blends realized retrospective outcomes with a modeled future-risk construct and a regional exogenous condition. This heterogeneity is deliberate. The aim is to capture both where graduates currently stand in the labor market and how exposed that position may be to emerging pressures.

Four percentile-ranked components receive equal weights:

Completion rate. Six-year graduation rate for first-time, full-time bachelor's degree-seeking students. Functions as both an outcome measure and a market position indicator.

Earnings-to-debt ratio. Median earnings at 10 years post-entry divided by median graduate debt. The primary return-on-investment measure.

Entry-level AI exposure (inverted). An institutional-level score derived from O*NET task-level AI exposure analysis, weighted for entry-level career pathways, crosswalked to degree fields via the NCES CIP-SOC mapping, and aggregated by institutional program mix. Inverted so low exposure equals high alignment. Because program diversification compresses institutional-level variation, this component contributes less between-institution discrimination than the other three; it is retained because it captures a qualitatively different risk dimension not reflected in any existing institutional health measure. Full methodology in Section 4.

Regional demographic trajectory. Projected percent change in high school graduates, 2024–2030, by state (WICHE).

Weighting and Scaling Rationale

All components receive equal weight. This is a deliberate choice, not an empirical claim that each contributes equally to institutional outcomes. The composite index literature (Saisana & Saltelli, 2011) documents that weighting is the single most consequential and least empirically grounded decision in any composite measure. I retain equal weighting because: (a) it avoids data-driven weights that would reflect historical relationships and thereby underweight emerging pressures; (b) it makes the framework interpretable for non-technical audiences; and (c) the sensitivity analysis provides a direct assessment of how weighting choices affect results. In practice, the AI exposure component contributes the least between-institution variance of any component; its influence on quadrant assignments is correspondingly modest despite its equal nominal weight.

Components are scaled using percentile ranks rather than standardization (z-scores). Percentile ranking compresses the distance between extreme values, a deliberate choice given that several components, particularly endowment per student, have extreme right skew that would dominate a z-score-based composite. The trade-off is that percentile ranking treats the distance between the 95th and 99th percentile as equivalent to the distance between the 50th and 54th. I report a z-score-based specification in the sensitivity analysis to confirm that the main findings are robust to this scaling choice.

Quadrant Assignment

Quadrant boundaries are set at the sample median on each axis, producing four categories: **High Capacity** (above both medians), **High Stress** (below both), **Market Misaligned** (above resilience, below market position), and **Structurally Exposed** (below resilience, above market position).

These labels describe structural position on specific indicators. They are evaluative in the sense that scoring higher reflects stronger positioning on the measures I track, but the indicators are analytical choices, not comprehensive assessments of institutional value or mission quality. An institution serving a vital access mission in a declining-enrollment region may score low precisely because it is doing important work in a structurally challenging context. The sensitivity analysis demonstrates that 69% of institutions change quadrant under at least one alternative specification; users should attend to both quadrant placement and stability scores.

4. AI Exposure Methodology

Why Institutions Need This Measure

Existing tools for assessing institutional health (financial composite scores, enrollment projections, closure prediction models) capture whether an institution can survive. They do not capture whether the career pathways its graduates enter will remain viable as AI reshapes labor markets. Yet this question has direct implications for institutional strategy: if the degree programs an institution is expanding to improve its competitive position are also the programs whose graduates face the highest AI task exposure, then the strategy that improves short-term market position may create longer-term vulnerability. The AI exposure component is designed to surface this dynamic.

The measure captures task-level *exposure* to AI capability, meaning whether AI systems can perform the tasks an occupation involves, not whether that exposure translates into substitution, augmentation, or occupational restructuring. I use the term "exposure" rather than "vulnerability" for this reason, though the framework's strategic logic assumes that high exposure warrants institutional attention regardless of which dynamic dominates.

Construction Pipeline

For each of 894 O*NET occupations, I computed an AI exposure score from two complementary approaches. The first classified 41 standardized work activities into AI-positive (routine cognitive: processing information, analyzing data, documenting, evaluating compliance) and AI-negative (physical, interpersonal: handling objects, assisting others, operating vehicles). The second scored individual tasks using the O*NET task-to-DWA hierarchy. The two approaches correlate at $\rho = 0.835$; the blended average is used.

AI disruption operates disproportionately at the entry level (Massenkoff & McCrory, 2026). The analysis weights occupation-level exposure by O*NET Job Zone: Zones 2–3 (typical entry points for college graduates) receive full weight; Zone 4 receives half weight; Zones 1 and 5 receive minimal weight. This weighting surfaces occupations where the first job on the career ladder is being restructured.

Occupation scores are mapped to degree fields via the NCES CIP-SOC crosswalk and aggregated to institutions using program mix weights. Full technical details, including the Job Zone weighting table and crosswalk diagnostics, are in Supplementary Appendix D.

Theoretical vs. Observed Exposure

I compared the theoretically-derived scores with observed real-world AI adoption rates from the Anthropic Economic Index. The field-level correlation is approximately zero ($\rho \approx -0.09$).

This divergence is consistent with a substantive interpretation: what AI *can* theoretically automate (driven by task characteristics) and what AI *is currently* automating (driven by user behavior, tool availability, and enterprise integration) are largely independent. Current adoption is dominated by coding and writing tasks; theoretical exposure is highest in business administration, legal support, and financial services. However, alternative interpretations (measurement artifacts, vendor-specific bias in single-platform adoption data, uneven regulatory constraints) are comparably plausible. The data cannot empirically distinguish among them.

If the substantive interpretation is correct, institutions may face two distinct exposure timelines. The near-term timeline is defined by current adoption patterns, which appear primarily augmentative. The medium-term timeline is defined by theoretical task exposure in fields where adoption has been slow but task susceptibility is high. Regulatory permeability likely moderates the pace of convergence: software development faces essentially zero regulatory friction to AI adoption; legal and financial services face compliance constraints that slow integration regardless of task suitability.

5. Results

Institutional Positioning

The quadrant distribution is approximately balanced by construction: 499 High Capacity, 491 High Stress, 277 Market Misaligned, and 289 Structurally Exposed. The distribution by Carnegie classification reveals pronounced stratification:

Carnegie Tier	High Stress	Structurally Exposed	Market Misaligned	High Capacity
R1	8 (5%)	12 (7%)	11 (6%)	146 (82%)
R2	21 (17%)	24 (20%)	25 (21%)	51 (42%)
RCU	46 (26%)	34 (19%)	36 (20%)	64 (36%)
Doctorate	57 (30%)	60 (32%)	23 (12%)	47 (25%)
Master's L/M	52 (46%)	14 (12%)	26 (23%)	21 (19%)
Master's S	62 (43%)	33 (23%)	29 (20%)	19 (13%)
Baccalaureate	150 (50%)	42 (14%)	63 (21%)	48 (16%)
Assoc/Bacc	23 (41%)	15 (27%)	11 (20%)	7 (12%)
Special Focus	72 (26%)	54 (20%)	53 (19%)	96 (35%)

Table 1. Distribution by Carnegie 2025 tier and quadrant. Cell values are counts with row percentages.

R1 institutions are structurally overdetermined: 82% fall in High Capacity. The Master's tiers show the opposite pattern, with 46% (L/M) and 43% (S) in High Stress. These are the tiers where institutional choices most determine outcomes. Baccalaureate institutions are the most concentrated in High Stress (50%). The RCU tier, new to the 2025 Carnegie system, distributes most broadly across quadrants, reflecting the heterogeneity of institutions with moderate research activity. Special Focus institutions show a bimodal pattern, with concentrations in both High Capacity (35%) and High Stress (26%).

[Figure 1 about here]

[Figure 2 about here]

The Bifurcation Pattern

The most striking finding is the enrollment × resilience cross-tabulation: 37.3% of institutions (581) are simultaneously losing enrollment and scoring below the median on resilience. Because enrollment trajectory is itself a resilience component, the co-occurrence is partly mechanical; however, the pattern holds under the specification that removes enrollment trajectory from the resilience composite (81.7% quadrant agreement), confirming that it reflects a genuine structural clustering rather than a definitional artifact. About 37.5% (584) are both growing and resilient. A further 12.6% (196) are growing from a weak institutional base, often through online or adult-learner expansion without corresponding improvements in completion or financial position. This pattern tracks with Grawe's (2018, 2021) demonstration that elite and highly selective institutions face fundamentally different demand conditions than the rest of the sector.

[Figure 3 about here]

AI Exposure Findings

Institutional-level AI exposure scores show modest variation (SD = 0.030) because program diversification smooths field-level differences. This compression is a predictable consequence of aggregation: most institutions offer programs across multiple fields, and the averaging attenuates field-level variation. The discriminating power of the measure is therefore strongest at the field level within institutions, not between institutions, which is why the field-level heatmap rather than the institutional index is the primary diagnostic tool for AI exposure.

The PSEO earnings comparison (572 institutions) shows a positive correlation between AI exposure and current earnings ($\rho = 0.257$, $p < 0.0001$). This is mechanically unsurprising: the occupations with the highest routine cognitive task content (engineering technology, computer science, business) currently command above-average entry-level wages precisely because complex information processing is economically valuable. The fields with low AI exposure (education, visual arts, biological sciences) involve tasks compensated differently in the current market. The strategic implication is that if disruption arrives at scale, it would affect the career pathways that currently justify the highest tuition and debt levels, creating a compounding problem for institutions that have expanded professional programs to improve their competitive position.

[Figure 4 about here]

Sensitivity and Robustness

I tested 18 alternative specifications plus a z-score scaling alternative against the baseline. The resilience axis is highly robust (all Spearman $\rho > 0.93$ across the original specifications). The post-college market position axis is more sensitive, particularly to AI specification (ρ ranges from 0.68 to 0.97). The AI exposure weight is the single most consequential parameter: at 50% weighting, 29% of institutions change quadrant; at 10%, only 12% change.

Key robustness checks: removing selectivity produces 85.7% quadrant agreement ($\rho = 0.938$), reassuring but not trivial; 154 institutions change, predominantly broad-access publics. Removing completion rate produces 87.3% agreement. Discounting sponsored research revenue produces 91.8% agreement. Replacing endowment level with endowment yield is more disruptive (77.2% agreement), reflecting the near-zero correlation between endowment size and single-year performance. The z-score specification produces 92.1% quadrant agreement with the baseline ($\rho = 0.987$ on resilience, $\rho = 0.985$ on market position), confirming that the main tier-level patterns are robust to the scaling choice. The 123 institutions that shift are distributed across all tiers and all quadrant transitions, with no tier showing agreement below 90%; R1 institutions show the highest agreement (96.0%), consistent with their overdetermined structural position.

Across all specifications, 488 institutions (31.4%) never change quadrant: 216 consistently High Capacity, 167 consistently High Stress. These stable positions are overdetermined and represent robust structural assessments. The remaining 69% are boundary cases whose classification depends on analytical assumptions.

Factor analysis (PCA on the 1,262 complete cases) yields a three-factor structure: Credential Outcomes (completion rate, earnings-to-debt; 29.1% variance), Institutional Character (revenue diversification, endowment, AI exposure; 18.9%), and Demand Environment (demographic trajectory, enrollment trend; 14.3%). The data's covariance structure does not naturally organize itself into two dimensions, and the dominant latent factor is best characterized as institutional hierarchy, the familiar gradient from well-resourced selective institutions to tuition-dependent open-access institutions. This result is expected given the framework's design intent. Factor analysis groups variables by historical covariance, which means it will always underweight emerging pressures that have not yet manifested in the outcome data. Demographic trajectory and enrollment momentum have uniqueness exceeding 0.87 in two-factor solutions precisely because they are orthogonal to the hierarchy signal, and it is that orthogonality that makes them diagnostically valuable. The 2×2 framework deliberately elevates these forward-looking indicators beyond what the covariance structure alone would support. Whether they prove predictive of future institutional outcomes is the framework's central empirical bet, and Kelchen et al. (2025) provide initial validation: they independently found that recent enrollment changes significantly improve prediction of closure and financial distress. Full factor analysis results across eight specifications are in Supplementary Appendix C.

Carnegie tier alone explains 25.1% of the variance in combined scores, and a Carnegie-modal baseline prediction achieves 45.3% quadrant accuracy, barely better than chance for a four-category outcome. The framework reclassifies 55% of institutions relative to this baseline. Its value lies in within-tier discrimination, particularly in the middle tiers where Carnegie classification tells you relatively little about an institution's actual structural position.

6. Discussion

This analysis extends the institutional financial health literature (Kelchen et al., 2025; Zemsky et al., 2020) by offering a framework that incorporates dimensions beyond financial solvency. The 277 Market Misaligned institutions are not at risk of imminent closure by any financial metric, yet their graduates enter career pathways that may be contracting or restructuring. Whether this constitutes meaningful institutional vulnerability or reflects different missions (liberal arts education, public-service preparation, regional embeddedness) is a question the framework surfaces but cannot resolve.

The near-zero correlation between theoretical and observed AI exposure is the analysis's most distinctive empirical pattern. If it reflects a genuine distinction between what AI can theoretically do and what it is currently being used for, then institutions may face two distinct exposure timelines rather than one, and the institutions most affected may not be those currently experiencing the highest AI adoption in their graduates' fields. They may instead be institutions whose graduates enter fields where the task structure is susceptible but adoption has been slow, the "latent vulnerability" zone the framework identifies. The positive correlation between AI exposure and current earnings adds urgency: disruption at scale would affect the career pathways that currently justify the highest tuition and debt levels. Institutions that have expanded professional and technical programs to improve their market position may find that the same programs create exposure to disruption those improvements do not yet reflect.

As discussed in the results, the factor analysis confirms that the eight components are dominated by institutional hierarchy. The framework's contribution is not in discovering new latent structure but in deliberately elevating indicators that the covariance structure undervalues. A purely data-driven composite would assign demographic trajectory and enrollment momentum minimal weight because they are largely orthogonal to the hierarchy signal. The 2×2 framework treats that orthogonality as informative rather than noisy, on the theory that the pressures most worth monitoring are those not already captured by the indicators institutions already track.

Limitations

The most consequential limitations are: (1) the AI exposure measure captures structural task characteristics but not institutional curricular response, so an institution actively integrating AI into its pedagogy may be better positioned than its exposure score suggests; (2) Scorecard earnings data covers only Title IV recipients, systematically understating outcomes at wealthier institutions; (3) the CIP-SOC crosswalk applies typical career pathways regardless of institutional prestige or selectivity; (4) the analysis is cross-sectional; and (5) institutions with fundamentally different funding models (military academies, tribal colleges) may produce technically accurate but strategically misleading scores under general-purpose indicators.

7. Implications for Practice and Future Research

For Institutional Leaders

The framework provides diagnostic visibility at the component level. An institution's quadrant placement is a starting point, but the component scores reveal which dimensions are driving positioning and which are most amenable to strategic intervention. Completion rates, program mix, and enrollment strategy are more within institutional control than endowment or state demographics; institutions near quadrant boundaries should examine whether targeted improvements on controllable dimensions could shift their structural position. The stability score, which reflects how many alternative specifications produce the same quadrant assignment, indicates whether an institution's classification is robust or contingent on analytical assumptions. The interactive tool at https://kylesaunders.com/university-map allows institutional researchers to locate their institution and examine component-level scores directly.

For State Systems and Coordinating Boards

The tier-level stratification pattern has resource allocation implications. The concentration of structurally stressed institutions in the Midwest and Northeast, disproportionately in the Baccalaureate and Master's tiers, suggests that closures and financial exigency declarations may cluster geographically. Many of these institutions function as anchor institutions in their communities (Harris & Holley, 2016), and where closures occur without adequate teach-out agreements, students face disrupted degree pathways and taxpayers absorb the cost of closed-school loan discharges (Kelchen et al., 2025). The framework provides an empirical basis for strategic conversations about resource allocation across systems that include institutions in multiple quadrants.

For Federal Policy

Three policy dimensions intersect with the framework's findings. First, the proposed elimination of the NSF Social and Behavioral Sciences directorate in the FY 2027 budget request, even if unlikely to proceed in its current form, signals a policy environment in which approximately $150 million in annual research funding cannot be taken for granted. The institutions most dependent on this funding are disproportionately in the framework's vulnerable quadrants: doctorate-granting institutions have the second-lowest revenue diversification (0.377) of any tier and 62% fall in the lower two quadrants. Second, the College Scorecard, which supplies several framework components, does not currently track AI-related labor market exposure or workforce restructuring, limiting its utility for forward-looking institutional assessment. Third, economists surveyed by Karger et al. (2026) favor worker retraining (71.8%) over universal basic income (37.4%) as the primary policy response to AI disruption — though these figures represent median expert forecasts with wide uncertainty bands — yet the institutions best positioned to deliver retraining at scale are the same regional institutions the framework identifies as structurally stressed.

For Future Research

Three extensions would strengthen the framework. First, longitudinal tracking as IPEDS and Scorecard data update annually, testing whether current positions predict subsequent institutional outcomes (closures, mergers, financial exigency declarations, enrollment declines, successful adaptation). Second, integration of PSEO actual employment destination data (available for 572 institutions) to replace crosswalk-based career pathway estimates with observed labor market outcomes. Third, development of an AI curriculum integration measure to complement exposure with adaptive capacity, transforming the framework from a one-dimensional risk measure into a risk-and-readiness assessment.

8. Conclusion

American higher education's structural divide is not a prediction; it is a current condition documented in publicly available data. The 37% of institutions simultaneously losing enrollment and scoring below the median on resilience, concentrated in baccalaureate colleges and smaller master's programs, face overlapping pressures from demographic contraction, revenue constraints, and potential AI-driven labor market restructuring. That the degree fields most theoretically exposed to AI disruption are currently the highest-earning fields sharpens the challenge: the programs many institutions have expanded to improve their competitive position may be creating exposure to the next wave of disruption. Whether the two-timelines pattern, the divergence between theoretical AI task exposure and observed adoption, persists or converges will shape which institutions face the most acute strategic challenges in the coming decade. The framework, dataset, and interactive tool are provided as public resources to support the institutional leaders, policymakers, and researchers working to navigate these challenges.

9. AI Use Disclosure

During the preparation of this work the author used Anthropic's Claude to assist with data assembly, statistical analysis, visualization, and sensitivity testing. The collaboration was iterative, with methodological decisions made by the author and implemented computationally by the AI system. All analytical decisions, theoretical framing, and interpretive judgments are the sole responsibility of the author, who reviewed and edited all outputs.

10. Data Access and Replication

All data sources are publicly available. The complete institutional dataset, AI exposure scores, and replication code are available at https://kylesaunders.com/university-map (Data & Downloads tab). An interactive tool for exploring institutional positions is available at the same URL; the tool has received more than 200,000 page views since its launch. The tool also provides an informal test of the framework's facial validity: users can look up institutions they know well and assess whether the component scores and quadrant placements accord with their own knowledge of those institutions' circumstances.

References

Bastedo, M. N. (2012). Organizing higher education: A manifesto. In M. N. Bastedo (Ed.), *The Organization of Higher Education* (pp. 3–17). Johns Hopkins University Press.

Bransberger, P., Falkenstern, C., & Lane, P. (2024). *Knocking at the College Door: Projections of High School Graduates* (11th ed.). WICHE.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. *arXiv:2303.10130*.

Felten, E., Raj, M., & Seamans, R. (2023). Occupational heterogeneity in exposure to generative AI. *SSRN 4414065*.

Galloway, S. (2020). USS University. *No Mercy / No Malice*.

Grawe, N. D. (2018). *Demographics and the Demand for Higher Education*. Johns Hopkins University Press.

Grawe, N. D. (2021). *The Agile College: How Institutions Successfully Navigate Demographic Changes*. Johns Hopkins University Press.

Harris, M. S., & Holley, K. (2016). Universities as anchor institutions. In M. B. Paulsen (Ed.), *Higher Education: Handbook of Theory and Research* (Vol. 31, pp. 393–439). Springer.

Hillman, N. W. (2014). College on credit: A multilevel analysis of student loan default. *The Review of Higher Education*, 37(2), 169–195.

Hillman, N. W. (2016). Geography of college opportunity. *American Educational Research Journal*, 53(4), 987–1021.

Jaquette, O., & Curs, B. R. (2015). Creating the out-of-state university. *The Journal of Higher Education*, 86(1), 47–77.

Karger, E., Kuusela, A., Abaluck, J., Tetlock, P., et al. (2026). Forecasting the economic effects of AI. *NBER Working Paper 35046*.

Kelchen, R., Ritter, D., & Webber, D. (2025). Predicting college closures and financial distress. *NBER Working Paper 33216*.

Ma, J., & Pender, M. (2024). *Trends in College Pricing and Student Aid 2024*. College Board.

Massenkoff, M., & McCrory, P. (2026). Labor market impacts of AI. *Anthropic Research*.

Morphew, C. C., & Huisman, J. (2002). Using institutional diversity to create dialogue for change. *The Review of Higher Education*, 26(3), 349–379.

National Association of College and University Business Officers. (2023). *2022 NACUBO Tuition Discounting Study*.

National Center for Education Statistics. (2023). *Digest of Education Statistics, 2023*.

National Student Clearinghouse Research Center. (2024). *Current Term Enrollment Estimates*.

Saisana, M., & Saltelli, A. (2011). Rankings and ratings: Instructions for use. *Hague Journal on the Rule of Law*, 3(2), 247–268.

Slaughter, S., & Rhoades, G. (2004). *Academic Capitalism and the New Economy*. Johns Hopkins University Press.

Tandberg, D. A. (2010). Politics, interest groups and state funding of public higher education. *Research in Higher Education*, 51(5), 416–450.

Tandberg, D. A. (2013). The conditioning role of state higher education governance structures. *The Journal of Higher Education*, 84(4), 506–543.

Third Way. (2023). The state of American higher education outcomes in 2023. *Third Way Report*.

U.S. Census Bureau. (n.d.). Post-Secondary Employment Outcomes (PSEO). Retrieved March 2026.

Zemsky, R., Shaman, S., & Baldridge, S. C. (2020). *The College Stress Test*. Johns Hopkins University Press.

Declarations

Conflict of Interest. The author declares no conflicts of interest. This research received no external funding.

Data Availability. All data are publicly available from IPEDS, College Scorecard, WICHE, O*NET, NCES CIP-SOC crosswalk, Census PSEO, and the Anthropic Economic Index. The complete dataset and replication code are at https://kylesaunders.com/university-map.

Figure 1

The Institutional Landscape: Resilience × Post-College Market Position

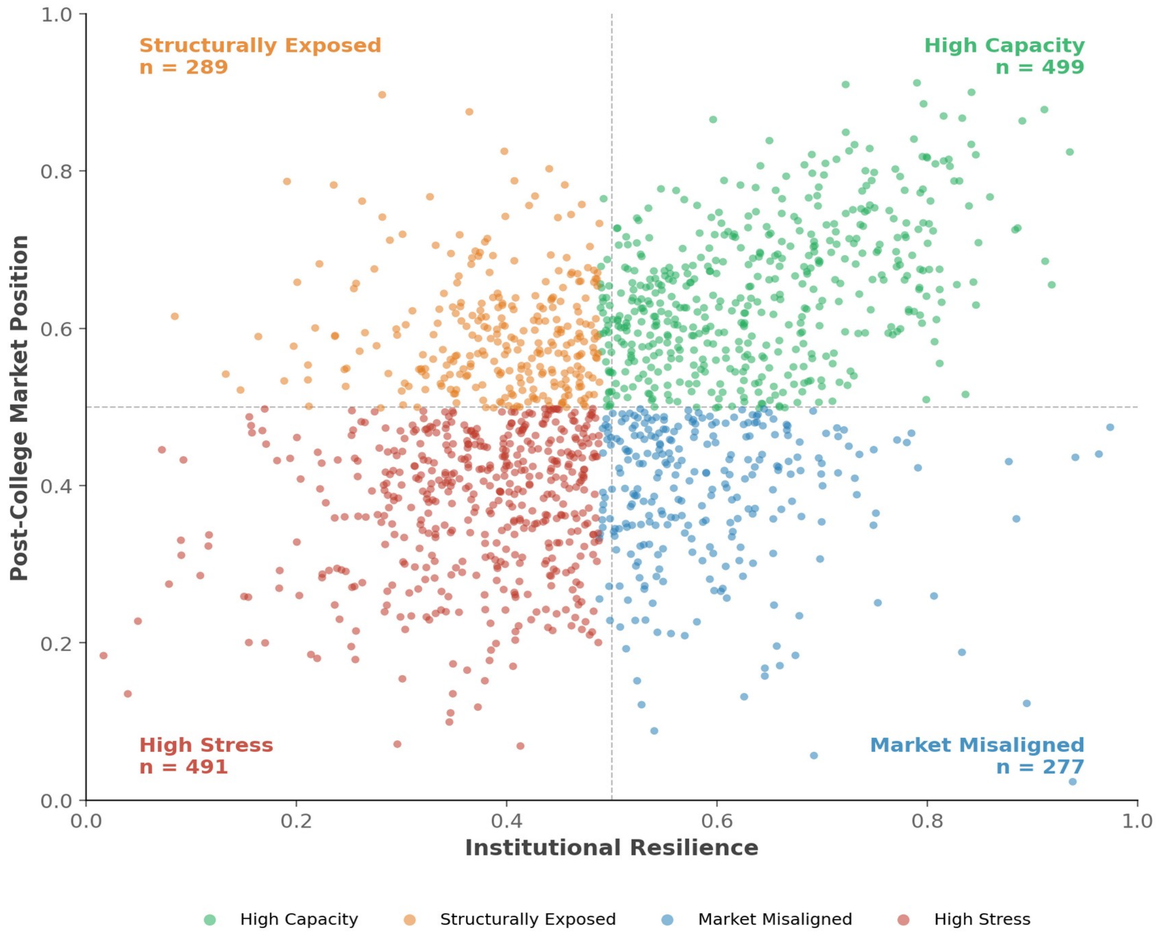

Note. Each point represents one of 1,556 four-year institutions. Quadrant boundaries are set at the sample median on each axis. Color indicates quadrant assignment.

Figure 2

Quadrant Distribution by Carnegie 2025 Tier

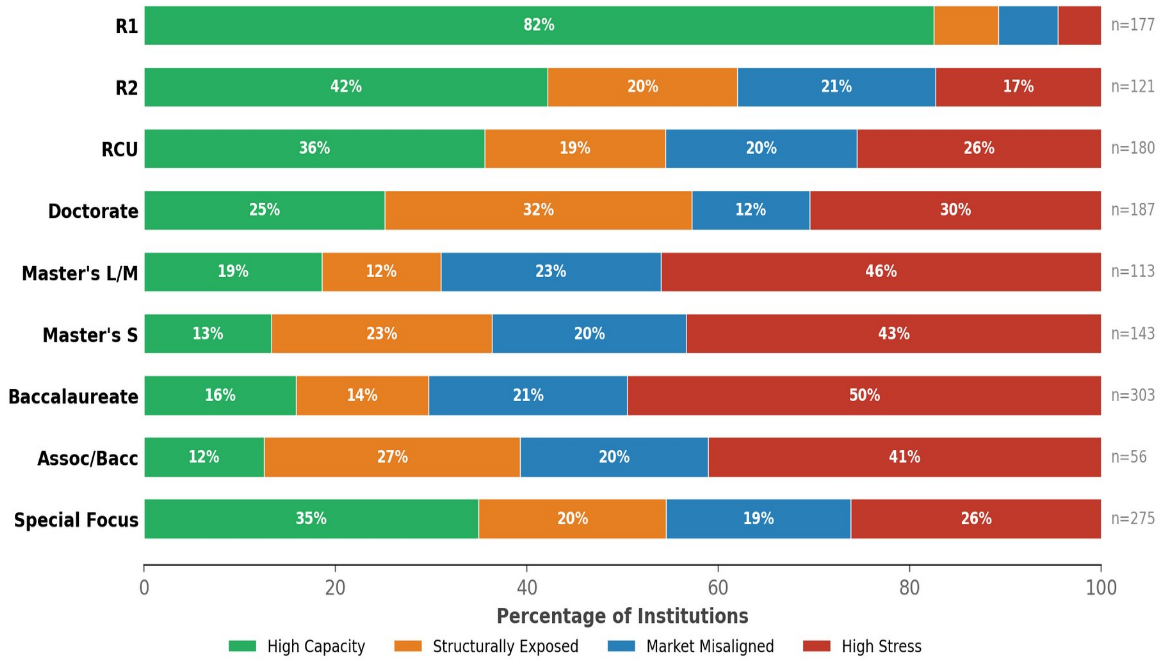

Note. Horizontal stacked bars show the percentage of institutions in each quadrant within each Carnegie tier. N for each tier shown at right.

Figure 3

Distribution of Combined Scores by Carnegie 2025 Tier

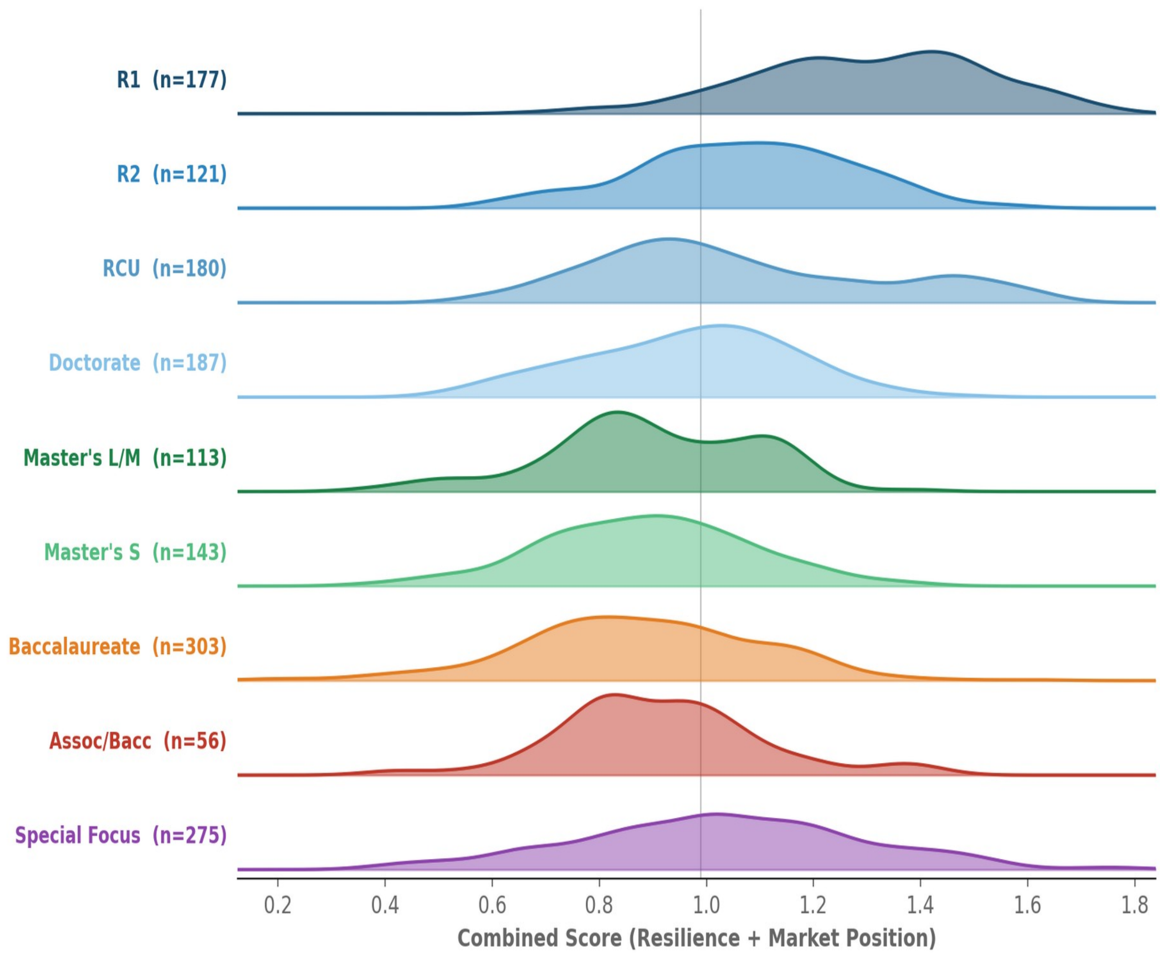

Note. Ridge plot showing the distribution of combined scores (Resilience + Post-College Market Position) by Carnegie tier. Vertical line marks the overall median.

Figure 4

Theoretical AI Exposure vs. Observed Adoption by Degree Field

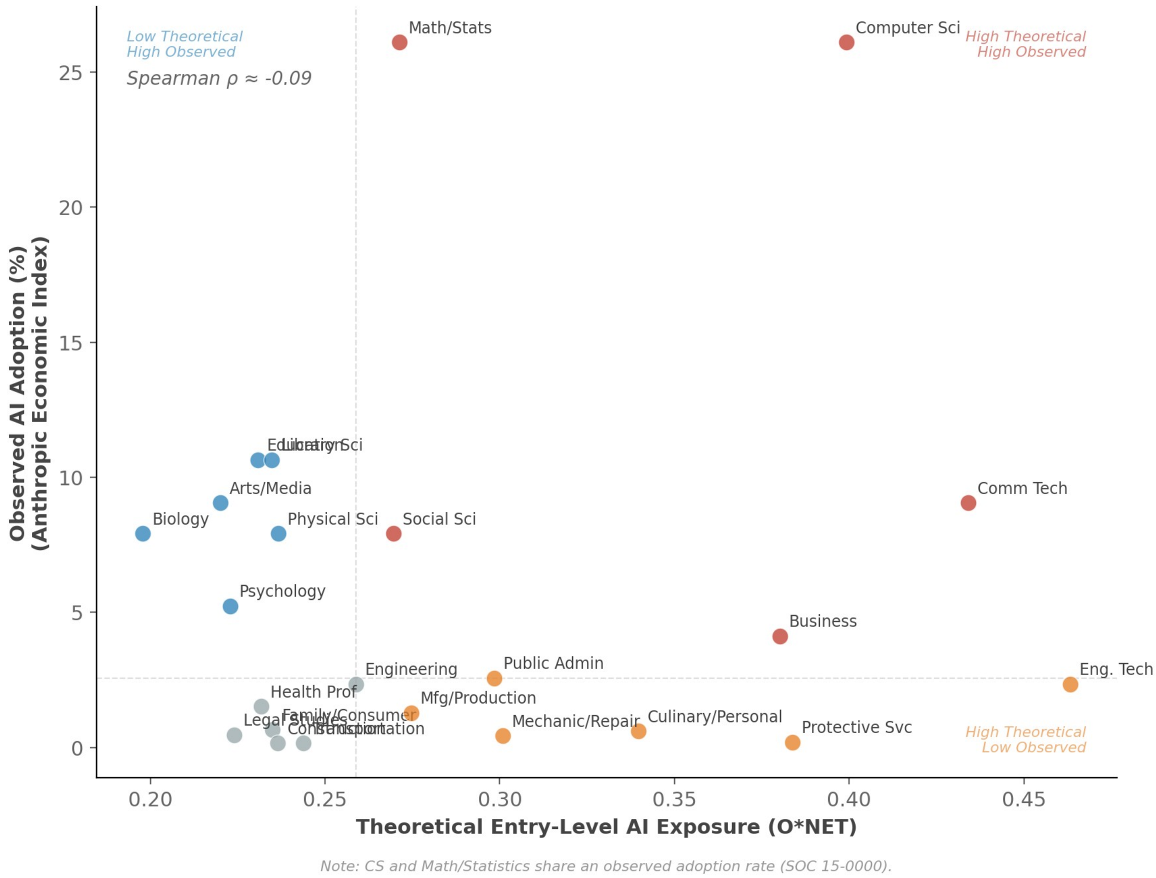

Note. Each point represents a 2-digit CIP degree field. Horizontal axis: theoretical entry-level AI task exposure (O*NET). Vertical axis: observed AI adoption rate (Anthropic Economic Index). Spearman ρ ≈ −0.09.