Supplementary Materials

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

v1.4 — April 2026

Appendix A: Data Sources and Variable Definitions

A.1 Data Sources

Source	Variables Used	Year	Access
IPEDS HD2024	Institution name, state, Carnegie 2025 classification (IC, RAD, SAEC), control type	2024	nces.ed.gov/ipeds
IPEDS F1A (GASB)	Total core revenues, net tuition revenue, endowment (public institutions)	FY2023 -24	nces.ed.gov/ipeds
IPEDS F2 (FASB)	Total revenues, net tuition revenue, endowment (private institutions)	FY2023 -24	nces.ed.gov/ipeds
IPEDS EFFY 2024 & 2019	12-month enrollment headcount	2024, 2019	nces.ed.gov/ipeds
IPEDS C2024_A	Completions by institution, CIP code, and award level	2024	nces.ed.gov/ipeds
College Scorecard	Admission rate, enrollment, completion rate, median debt, net price, program mix (PCIP)	Most recent cohort	collegescorecard.ed.gov
WICHE (11th ed.)	High school graduate projections by state, 2024–2041	Dec 2024	wiche.edu/knocking
O*NET Database	Work activities, task statements, task ratings, job zones, occupation data	v29.0	onetcenter.org
Anthropic Economic Index	Observed AI task usage from Claude conversations	Aug 2025	huggingface.co/ datasets/Anthropic/ EconomicIndex
NCES CIP-SOC Crosswalk	CIP 2020 to SOC 2018 occupation mappings	2020	nces.ed.gov
Census PSEO	Median earnings at 1, 5, 10 years by institution × CIP	Various cohorts	lehd.ces.census.gov/ data/pseo
Carnegie/ACE 2025 Public Data File	Institutional Classification (IC2025), Award Level Focus, Size, Setting, Highest Degree	Oct 2025	carnegieclassifications. acenet.edu
Carnegie/ACE 2025 RAD Public Data File	Research Activity Designation (R1/R2/RCU), HERD expenditures, research doctorate counts	Oct 2025	carnegieclassifications. acenet.edu
Carnegie/ACE 2025 SAEC Public	Student Access and Earnings Classification, Pell ratios, URM ratios, earnings data	Oct 2025	carnegieclassifications. acenet.edu

Data File

A.2 Selection Criteria

Institutions included if all conditions met from IPEDS HD2024:

• ICLEVEL = 1 (4-year institution)

• DEGGRANT = 1 (degree-granting)

• CONTROL ∈ {1, 2} (public or private not-for-profit)

• Four-year, degree-granting institutions classified under the 2025 Carnegie system (Institutional Classification codes 1–31; Research Activity Designation where applicable). The legacy variable C21BASIC ∈ {15–23} is retained for backward compatibility.

• CURROPER = 1 (currently operating, per College Scorecard)

A.3 Carnegie 2025 Classification and Tier Derivation

The 2025 Carnegie Classification system (as reported in IPEDS HD2024) replaced the single Basic Classification with three separate systems: the Institutional Classification (31 categories based on award level focus, academic program mix, and size), the Research Activity Designation (R1, R2, and the new Research Colleges and Universities category), and the Student Access and Earnings Classification. I derive a 9-tier grouping variable for stratified analyses using the Research Activity Designation as the primary axis for research institutions and Award Level Focus for non-research institutions:

Tier Code	Tier Label	Derivation Rule	N in sample

1	R1	Research Activity Designation = 1	177
2	R2	Research Activity Designation = 2	124
3	RCU	Research Activity Designation = 3	182
4	Doctorate	IC2025 ∈ {6,7,8,16,17,18} (Doctorate-granting, no research designation)	189
5	Master's L/M	IC2025 ∈ {9,19} (Master's Large/Medium)	113
6	Master's S	IC2025 ∈ {10,20} (Master's Small)	147
7	Baccalaureate	IC2025 ∈ {5,14,15} (Baccalaureate)	311
8	Assoc/ Bacc	IC2025 ∈ {1,2,3,4,11,12,13} (Associate/Baccalaureate mixed)	62
9	Special Focus	IC2025 ≥ 21 (Special Focus institutions)	303

The legacy 2021 Basic Classification variable (C21BASIC) and the raw 2025 codes (IC2025, RESEARCH2025) are included in the public dataset for researchers who prefer alternative groupings.

A.3a Carnegie 2025 Data File Integration

The 2025 Carnegie Classification was released by the Carnegie Foundation and the American Council on Education (ACE) in April 2025, with technical documentation and public data files updated October 10, 2025. IPEDS HD2024 incorporated the full 2025 classification suite as seven new variables alongside the legacy 2021 Basic Classification.

Three Carnegie/ACE public data files were used to construct the tier variable and to include supplementary classification fields in the dataset:

2025 Public Data File (3,929 institutions): Contains the Institutional Classification (IC2025, 31 categories), Award Level Focus, Academic Program Mix (undergraduate and graduate), Size, Setting, and Highest Degree for every classified institution. Also includes the legacy 2021 Basic Classification code. The IC2025 variable is the primary input for non-research tier assignment.

2025 RAD Public Data File (543 institutions): Contains the Research Activity Designation for institutions meeting research thresholds. Thresholds are: R1 (≥$50M aggregate HERD FY21–23 research expenditures and ≥70 research doctorates), R2 (≥$5M and ≥20 doctorates), RCU (≥$2.5M research expenditures). The RAD variable takes precedence in tier assignment — any institution with a research designation is classified by that designation regardless of its IC2025 code.

2025 SAEC Public Data File (3,928 institutions): Contains the Student Access and Earnings Classification, which groups institutions by Pell Grant recipient ratios, underrepresented minority ratios, and post-enrollment earnings. The SAEC is not used in the tier derivation but is available in the public dataset for researchers interested in access-oriented stratification.

Match rates: 1,608 of 1,609 institutions in the analytical universe matched to the Carnegie 2025 Public Data File by UNITID. The single unmatched institution (Northeastern University Oakland, UNITID 118888) is a recently established branch campus not yet classified.

A.4 Missingness Composition

Of 1,609 institutions meeting inclusion criteria, 53 (3.3%) lack sufficient data to compute both composite scores. The table below compares institutions with complete versus incomplete data.

Characteristic	Both Scores (n=1,556)	Insufficient Data (n=53)
Mean enrollment (UGDS)	5,920	1,290

Median enrollment	2,980	650
% Private nonprofit	73.1%	84.9%
% Special Focus	17.7%	52.8%
% Baccalaureate	19.5%	15.1%
% Master's (L/M + S)	16.5%	7.5%
% R1/R2	19.2%	5.7%
% Assoc/Bacc	3.6%	11.3%

Missingness is concentrated among Special Focus institutions (52.8% of unmapped institutions) and smaller private schools that do not report all College Scorecard outcome variables. The mapped universe therefore somewhat underrepresents the smallest and most resource-constrained institutions — precisely the population most likely to be structurally stressed. The most common missing variables are admission rate and debt-to-earnings metrics.

A.5 Comparison to Galloway (2020)

Galloway used approximately 436 institutions from the U.S. News & World Report ranked universe. The present universe is deliberately broader: U.S. News rankings are methodologically contested, several major institutions have withdrawn, and restricting to ranked institutions excludes many of the most vulnerable. I present the full universe with ~400 prominent institutions labeled.

Appendix B: GASB/FASB Finance Correction

IPEDS reports financial data for public institutions under Governmental Accounting Standards Board (GASB) standards (Form F1A) and for private nonprofit institutions under Financial Accounting Standards Board (FASB) standards (Form F2). The F1A form's revenue fields (F1A01, F1A05) are cumulative running totals within the survey instrument, not individual line items. This distinction is not well-documented in the IPEDS data dictionaries and represents a common source of error in cross-sector financial analyses.

Incorrect approach (initial analysis): Using F1A01 as total revenue and F1A05 as net tuition revenue produced tuition dependence ratios of ~1.0 for nearly all public institutions — an obviously incorrect result given that public institutions derive substantial revenue from state appropriations, grants, and contracts.

Corrected approach:

• Public institutions (GASB): F1D01 (total core revenues), F1B01 (net tuition and fees), F1H02 (endowment end-of-year market value)

• Private nonprofit institutions (FASB): F2A01 (total revenues), F2B01 (net tuition and fees), F2H02 (endowment end-of-year market value)

Effect on results:

Metric	Before correction	After correction
Public mean tuition dependence	98.6%	16.7%
Private NP mean tuition dependence	78.8%	77.8%
Public mean resilience score	0.403	0.510
Private NP mean resilience score	0.555	0.493
Publics in High Capacity quadrant	116	180

The correction eliminated a substantial systematic bias against public institutions on the resilience axis.

Appendix C: Framework Structure and Sensitivity Analysis

C.1 Eight-Component Framework Structure

The framework comprises eight components organized into two four-component scales:

Resilience (4 components, equal weight 0.25 each):

• Endowment per student (percentile rank)

• Revenue diversification (1 − net tuition revenue / total core revenues, percentile rank)

• Enrollment trajectory (5-year % change in 12-month headcount, 2019–2024, percentile rank)

• Selectivity (1 − admission rate, percentile rank)

Post-College Market Position (4 components, equal weight 0.25 each):

• Completion rate (6-year graduation rate, percentile rank)

• Earnings-to-debt ratio (10-year median earnings / median graduate debt, percentile rank)

• AI exposure (inverted, percentile rank)

• Demographic trajectory (WICHE projected % change in high school graduates, 2024–2030, by state)

Both scales are normalized to 0–1 percentile rank, then combined using equal weighting to form composite scores on the horizontal (Resilience) and vertical (Market Position) axes of the institutional typology.

C.2 Full Sensitivity Results

Specification	ρ(Resilience)	ρ(Market Pos.)	Same Quadrant

Selectivity removal	0.951	1.000	85.7%
Double-weight completion	0.989	1.000	91.8%
Adding earnings back	0.987	0.992	89.1%
Double-weight AI	1.000	0.968	88.8%
No selectivity (X-axis)	0.951	1.000	85.7%
No earn/debt (Y-axis)	1.000	0.911	84.0%
No demographics (Y-axis)	1.000	0.924	83.7%
No completion (both)	0.959	1.000	82.7%
No AI exposure (Y-axis)	1.000	0.918	82.4%
No rev diversification (X-axis)	0.980	1.000	82.2%
No enrollment (X-axis)	0.988	1.000	81.7%
No endowment (X-axis)	0.981	1.000	81.4%
Completion on resilience (old structure)	0.878	1.000	72.5%
RevDiv: ICR discount (33% HERD)	0.933	1.000	91.8%
RevDiv: Full HERD removal	0.932	1.000	91.5%
Endowment yield replaces level	0.733	1.000	77.2%
Half-weight endowment	0.969	1.000	92.5%
Admission yield replaces selectivity	0.789	1.000	80.2%

C.3 Institution-Level Stability

• 488 institutions (31.4%) never change quadrant across sensitivity specifications

- High Capacity: 216 (stably high-capacity) - High Stress: 167 (stably high-stress) - Structurally Exposed: 60 - Market Misaligned: 45

• 28 most unstable institutions change quadrant in 5–7 of specifications, concentrated near axis medians

C.4 Z-Score (Probit Transform) Robustness Check

The baseline framework uses percentile ranks (0–1), which compress the distance between extreme values. To test robustness to this scaling choice, each percentile-ranked component was transformed using the probit function (inverse normal CDF): values were ranked within non-missing cases, scaled to the open interval (0,1) using rank/(N+1), and transformed via Φ^{-1}. This decompresses the extremes that percentile ranking compressed, giving disproportionately more spread to institutions at the tails.

Metric	Value
Quadrant agreement with baseline	92.1% (1,433/1,556)
Spearman ρ (Resilience axis)	0.987
Spearman ρ (Market Position axis)	0.985
Institutions that shift quadrant	123

Per-tier agreement:

Carnegie Tier	Agreement
R1	170/177 (96.0%)
R2	109/121 (90.1%)
RCU	168/180 (93.3%)
Doctorate	172/187 (92.0%)
Master's L/M	103/113 (91.2%)
Master's S	132/143 (92.3%)
Baccalaureate	277/303 (91.4%)
Assoc/Bacc	52/56 (92.9%)
Special Focus	249/275 (90.5%)

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. The most common transitions are symmetric movements across axis medians (Structurally Exposed ↔ Market Misaligned: 20; High Capacity ↔ Market Misaligned: 18; High Capacity ↔ High Stress: 16). These results confirm that the main tier-level patterns and institutional rankings are robust to the scaling choice.

C.5 Note on Log Transform

The baseline model applies a log transform to endowment per student before percentile ranking. Because percentile ranking is a monotonic transform, log(x) and x produce identical percentile ranks. The log transform is analytically inert and is retained only for documentation transparency.

C.6 Revenue Diversification: Sponsored Research Discount (v1.3)

Revenue diversification (1 − net tuition / total core revenues) may overstate the true diversification of research-intensive institutions because a substantial share of non-tuition revenue consists of sponsored research funding that is restricted, project-specific, and not freely deployable against financial shocks. I test whether discounting sponsored research revenue changes institutional rankings.

Using HERD (Higher Education Research and Development) expenditure data from the Carnegie 2025 RAD file, I identify 467 institutions with both HERD and IPEDS finance data. At R1 institutions, HERD expenditures average 48.1% of total core revenues (median 21.0%); at R2 institutions, 9.9% (median 5.5%); at RCU institutions, 8.3% (median 3.4%).

I test two specifications: (a) removing the indirect cost recovery portion (33% of HERD, the approximate fraction representing discretionary overhead reimbursement) from total revenues before computing revenue diversification, and (b) removing all HERD expenditures from total revenues (an extreme upper bound). In both cases, all institutions are re-ranked together and quadrants are recomputed from scratch.

Specification	ρ(Resilience)	ρ(Market Pos.)	Same Quadrant
ICR discount (33% of HERD removed)	0.933	1.000	91.8%
Full HERD removal (100%)	0.932	1.000	91.5%

Both specifications produce high quadrant agreement with the baseline (>91%), and the difference between the ICR discount and the extreme full-removal specification is negligible (0.3 percentage points). Quadrant changes are concentrated in non-research tiers (Doctorate 13.8%, Master's L/M 11.5%) rather than in R1 (0.6%) or R2 (2.4%), because re-ranking research institutions' RevDiv downward pushes some non-research institutions above the median. The dominant transition pattern is symmetric movement across the resilience median: High Stress ↔ Market Misaligned and High Capacity ↔ Structurally Exposed.

The near-identical results under partial versus full HERD removal indicate that the framework's resilience rankings are robust to the treatment of sponsored research revenue. Research intensity does inflate raw diversification scores, but because the framework uses percentile ranks — and research institutions move together — the relative ordering is largely preserved.

C.7 Endowment Encumbrance and Admission Yield (v1.3)

Two related concerns motivate this sensitivity: (a) endowment levels may overstate institutional resilience because the vast majority of endowment assets are restricted by spending rules (typically 4–5% annual draw) and donor restrictions, and (b) selectivity (1 − admission rate) may not capture demand-side resilience as well as admission yield (enrolled / admitted), which reflects an institution's ability to convert offers into matriculants.

Endowment yield replacement. I replace endowment level (the baseline R_ENDOW component) with one-year endowment yield — (EOY − BOY) / BOY from IPEDS FY2023–24 finance data — as a proxy for endowment performance rather than endowment size. The Spearman correlation between endowment level and endowment yield is $\rho = -0.140$ ($N = 1,484$), confirming that yield captures fundamentally different information than level. This specification produces the largest quadrant disruption among all v1.3 sensitivities:

Specification	ρ(Resilience)	ρ(Market Pos.)	Same Quadrant
Endowment yield replaces level	0.733	1.000	77.2%
Half-weight endowment	0.969	1.000	92.5%
Admission yield replaces selectivity	0.789	1.000	80.2%

The endowment yield specification reshuffles 22.8% of institutions, with changes distributed broadly across tiers (Assoc/Bacc 32.3%, R2 and Master's L/M 27.4%, RCU 25.8%). This reflects the near-zero correlation between endowment size and single-year return — large endowments do not systematically outperform small ones in any given year. By contrast, when endowment is merely down-weighted (half-weight specification, compressing R_ENDOW toward 0.5), quadrant agreement remains high at 92.5%, indicating that the framework is not unduly dependent on endowment's specific weight.

Admission yield. Replacing selectivity with admission yield (ENRLT / ADMSSN from ADM2024, $N = 1,378$ reporting) produces 80.2% quadrant agreement. The Spearman correlation between baseline selectivity and admission yield is only $\rho = 0.169$ ($N = 1,371$) — the two measures capture largely independent demand signals. Changes are concentrated in RCU (26.4%), Master's L/M (23.0%), and Doctorate (22.2%) tiers, with the dominant transition being symmetric movement across the resilience median.

Interpretation. The endowment yield result is the most informative: it demonstrates that the baseline framework's resilience axis captures endowment *size* (a stable, slow-moving institutional characteristic) rather than endowment *performance* (a volatile, year-specific outcome). This is by design — the framework measures structural capacity, not short-term financial results — but users interested in financial performance should note the distinction. The half-weight specification confirms that the framework does not critically depend on endowment magnitude. The admission yield result suggests that replacing selectivity with yield would produce a meaningfully different resilience ranking, though neither measure is unambiguously superior as a demand indicator.

C.8 Factor Analysis

Sampling adequacy diagnostics (8 percentile-ranked components, N = 1,262):

The Kaiser-Meyer-Olkin (KMO) measure of sampling adequacy and Bartlett's test of sphericity were computed as prerequisites for factor analysis.

Diagnostic	Value
KMO overall	0.63 ("Mediocre")
Bartlett's $\chi^2(28)$	1900.58 ($p < .001$)
Correlation matrix determinant	0.221

The moderate KMO reflects the framework's deliberate inclusion of components that are empirically orthogonal to the dominant institutional hierarchy dimension. Demographic trajectory (MSA = 0.54) and revenue diversification (MSA = 0.56) have the lowest individual MSA values, consistent with their low correlations with other components. Selectivity has the highest MSA (0.82). Bartlett's test confirms that the correlation matrix differs significantly from an identity matrix and is factorable.

Per-variable Measures of Sampling Adequacy (MSA):

Component	MSA
Selectivity	0.822
AI exposure (inv)	0.738
Endowment/student	0.650
Completion rate	0.608
Enrollment trend	0.597
Earnings-to-debt ratio	0.571
Revenue diversification	0.564
Demographic trajectory	0.537

PCA eigenvalues (8 components):

Component	Eigenvalue	Variance Explained	Cumulative
PC1	2.33	29.1%	29.1%
PC2	1.51	18.9%	48.0%
PC3	1.15	14.3%	62.3%
PC4	0.82	10.3%	72.6%
PC5	0.79	9.9%	82.5%
PC6	0.67	8.4%	90.9%
PC7	0.46	5.7%	96.6%
PC8	0.27	3.4%	100.0%

Parallel analysis (Horn, 1965; 1,000 random-data replications):

Component	Observed Eigenvalue	Random 95th Percentile	Retain?
PC1	2.33	1.16	Yes
PC2	1.51	1.10	Yes
PC3	1.15	1.06	Yes
PC4	0.82	1.03	No

Both the Kaiser criterion (eigenvalue ≥ 1.0) and parallel analysis converge on a three-factor solution. The third eigenvalue (1.15) exceeds the 95th percentile of random data (1.06), confirming that the third factor captures more variance than expected by chance. Results are reported across eight specifications: 2 variable scalings (raw standardized, percentile rank) × 2 rotations (varimax, promax oblique) × 2 factor counts (2, 3).

Missingness note: The factor analysis uses 1,262 complete cases (of 1,556 mapped institutions). The 294 institutions excluded due to missing component data are disproportionately smaller private institutions and Special Focus schools: 42.5% of excluded institutions (where quadrant assignment is available) fall in High Stress, compared to the 31.5% rate among included institutions. The factor structure may not fully generalize to the most structurally stressed portion of the population.

Two-factor varimax-rotated loadings (percentile rank variables, 8 components, N=1,262):

Component	Factor 1: Credential Outcomes	Factor 2: Institutional Character	Communality (h^2)	Uniqueness (u^2)
Earnings/debt ratio	**−0.807**	−0.329	0.759	0.241

Completion rate	**−0.860**	+0.209	0.784	0.216
Revenue diversification	−0.124	**−0.800**	0.656	0.344
Endowment/ student	−0.554	**+0.596**	0.662	0.338
Selectivity	−0.507	+0.297	0.345	0.655
AI exposure (inverted)	−0.108	**+0.637**	0.418	0.582
Enrollment trend	−0.345	−0.087	0.126	**0.874**
Demographic trajectory	+0.284	+0.079	0.087	**0.913**

Two-factor varimax-rotated loadings (raw standardized variables, 7 components):

Note: Enrollment trend is available only as a percentile rank derived from the 2019–2024 enrollment change. Because its raw distribution has strong ceiling and floor effects (many institutions near zero change), the raw standardized specification omits it. The percentile rank specifications above include all 8 components.

Component	Factor 1: Quality + Wealth	Factor 2: Program Character	Uniqueness
Earnings/debt ratio	**−0.879**	+0.182	0.194
Completion rate	**−0.795**	−0.195	0.336
Endowment/student	**−0.699**	−0.123	0.501
Selectivity	**−0.670**	−0.241	0.491
Revenue diversification	−0.024	**+0.791**	0.375
AI exposure (inverted)	−0.022	**−0.758**	0.426
Demographic trajectory	+0.161	−0.053	**0.973**

Note: In raw standardized data, endowment loads on Factor 1 with outcomes rather than on Factor 2, because the extreme right skew of the raw endowment distribution creates a stronger endowment-outcomes correlation. The percentile rank transformation compresses this skew, allowing endowment to separate onto its own dimension. Both results are substantively interpretable; the core finding — that forward-looking indicators have near-zero loadings — is identical.

Three-factor promax (oblique) loadings (percentile rank variables, 8 components, N=1,262):

Component	F1: Credential Outcomes	F2: Institutional Character	F3: Demand Environment	Communality (h^2)	Uniqueness (u^2)
Earnings/debt ratio	**−0.850**	−0.389	+0.111	0.769	0.231
Completion rate	**−0.826**	+0.139	+0.081	0.790	0.210
Revenue diversification	−0.227	**−0.807**	−0.003	0.656	0.344
Endowment/ student	−0.461	**+0.543**	−0.011	0.662	0.338
Selectivity	−0.413	+0.229	−0.288	0.427	0.573
AI exposure (inverted)	−0.042	**+0.635**	+0.139	0.439	0.561
Enrollment trend	−0.247	−0.167	**−0.671**	0.580	0.420
Demographic trajectory	+0.403	+0.046	**−0.760**	0.660	0.340

Factor correlations: F1-F2 = −0.221, F1-F3 = +0.175, F2-F3 = −0.097.

Three-factor promax (oblique) loadings (raw standardized variables, 7 components):

Component	F1: Credential Outcomes	F2: Program Character	F3: Momentum	Uniqueness

Earnings/debt ratio	**−0.894**	+0.233	−0.043	0.194
Completion rate	**−0.766**	−0.159	−0.103	0.329
Endowment/student	**−0.672**	−0.056	+0.109	0.484
Selectivity	**−0.632**	−0.163	+0.174	0.457
Revenue diversification	−0.118	**+0.797**	+0.039	0.370
AI exposure (inverted)	+0.068	**−0.764**	−0.046	0.413
Demographic trajectory	+0.176	+0.078	**+0.796**	0.361

Factor correlations: F1-F2 = +0.188, F1-F3 = −0.088, F2-F3 = −0.100.

Factor-implied weights vs. equal weights (communality-based, F1+F2):

Component	Equal Weight	Factor-Implied
Earnings/debt ratio	0.125	0.24
Completion rate	0.125	0.19
Revenue diversification	0.125	0.19
Endowment/student	0.125	0.14
AI exposure	0.125	0.11
Selectivity	0.125	0.06
Demographics	0.125	0.04
Enrollment trend	0.125	0.02

Note: Equal weights shown as 0.125 (1/8) for cross-component comparison. Within each axis, each component receives 0.25. Factor-implied weights are derived from the F1+F2 communality of each component, normalized to sum to 1.0. The data's covariance structure allocates most variance to institutional hierarchy indicators (earnings, completion, revenue diversification) and minimal variance to forward-looking indicators (enrollment trend, demographics) — precisely the indicators the equal-weight framework deliberately elevates.

Equal-weight vs. factor-derived quadrant agreement: 61.3%.

Summary of robustness across specifications:

Feature	**Stable across all 8 specs?**
Earnings anchors F1 (loadings 0.81–0.89)	Yes
Completion crossloads F1 and F2	Yes
Forward-looking indicators invisible in 2F (uniqueness > 0.87)	Yes
Forward-looking indicators anchor F3 when extracted	Yes
Oblique rotation confirms moderate factor correlations	Yes
AI exposure loads on F2, not F1	Yes
Specific composition of F2	No — varies by scaling
Where endowment loads	No — F1 in raw, F2 in ranked

Interpretation: Results are robust to extraction method, rotation, and variable scaling on the central finding: the data does not naturally organize into "Resilience" and "Post-College Market Position" along equal-weight lines. The primary dimension is always Credential Outcomes; the secondary dimension captures institutional character whose specific composition varies by specification. Forward-looking indicators (demographic trajectory) are consistently invisible to the two-factor structure and emerge as a distinct "Momentum" factor only when three factors are extracted, with eigenvalues just above the Kaiser threshold. This third factor corresponds precisely to the indicators the equal-weight framework deliberately elevates — indicators whose effects have not yet had time to register in the backward-looking covariance structure. Oblique rotation confirms moderate interfactor correlations ($|r| = 0.10–0.22$) without altering the substantive interpretation.

Additional specification: Omitting earnings-to-debt ratio. Because the earnings-to-debt ratio is a composite of two variables already represented individually (median earnings and institutional debt levels that correlate with institutional type), it may mechanically inflate covariance between the outcomes and institutional-character dimensions. An 8-variable specification omitting this ratio produces a cleaner two-factor structure: Factor 1 (completion −0.869, earnings −0.881, endowment −0.538, selectivity −0.432) and Factor 2 (AI exposure +0.680, revenue diversification −0.783). The three-factor solution continues to produce a Momentum factor anchored by demographics (−0.691). This specification suggests that the eight components are dominated by a single latent dimension — institutional hierarchy — with AI exposure and demographics representing independent contextual pressures that do not correlate with hierarchy and therefore provide genuinely new information beyond what prestige or outcome quality alone can capture.

Appendix D: AI Exposure Detailed Methodology

D.1 Work Activity Classifications

AI-Positive (7 activities): Processing Information, Analyzing Data or Information, Working with Computers, Documenting/Recording Information, Evaluating Information to Determine Compliance with Standards, Performing Administrative Activities, Getting Information

AI-Negative (11 activities): Performing General Physical Activities, Handling and Moving Objects, Operating Vehicles/Mechanized Devices/Equipment, Controlling Machines and Processes, Repairing and Maintaining Mechanical Equipment, Repairing and Maintaining Electronic Equipment, Inspecting Equipment/Structures/Materials, Assisting and Caring for Others, Performing for or Working Directly with the Public, Coaching and Developing Others, Training and Teaching Others

Neutral (23 activities): All remaining work activities.

D.2 Entry-Level Weighting (Job Zone)

Job Zone	Description	Weight	Rationale
1	Little/no preparation	0.2	Pre-college labor
2	Some preparation	1.0	Entry-level for graduates
3	Medium preparation	1.0	Entry-level for graduates
4	Considerable preparation	0.5	Augmented, not replaced
5	Extensive preparation	0.2	Requires years of training

D.3 Task-Level AI Scoring

Individual tasks scored via O*NET task-to-DWA hierarchy: each task maps to Detailed Work Activities, which map to the 41 classified work activities. Task AI score = proportion of linked DWAs classified AI-positive minus AI-negative. Occupation-level score = importance-weighted average of task scores. Correlation between work-activity-based and task-based approaches: $\rho = 0.835$.

D.4 Face Validity

Highest AI exposure occupations: Medical Transcriptionists (1.000), Credit Authorizers (1.000), Statistical Assistants (0.996), Title Examiners (0.976), Paralegals (0.889)

Lowest AI exposure occupations: Electrical Power-Line Installers (0.000), Landscaping Workers (0.061), Terrazzo Workers (0.061), Paving Equipment Operators (0.083)

Highest entry-level-weighted exposure: Eligibility Interviewers (0.926), Title Examiners (0.906), Payroll Clerks (0.886), Billing Clerks (0.867), Paralegals (0.861)

D.5 Conceptual Scope

The AI exposure measure captures task-level *exposure* — whether AI systems can perform the tasks central to an occupation — not whether that exposure will manifest as task *substitution* (AI replacing workers), task *augmentation* (AI enhancing worker productivity), or occupational *restructuring* (workflow reorganization that changes the occupation's task composition). A high exposure score indicates structural susceptibility to AI involvement in the occupation's task content; it does not predict the labor market outcome. I use "exposure" rather than "vulnerability" or "risk" for this reason.

D.6 Tier 2: 6-Digit CIP Analysis

For 459 R1, R2, and Doctorate-level institutions, IPEDS C2024_A completions at the 6-digit CIP level provide finer granularity than the 2-digit PCIP shares used in Tier 1. The CIP-SOC crosswalk maps at the 6-digit level, with 4-digit and 2-digit fallbacks for unmatched codes. The institutional AI exposure range at Tier 2 (0.224–0.323) is compressed relative to Tier 1 because large research universities have diversified program mixes that smooth field-level variation.

D.7 Theoretical vs. Observed Exposure Comparison

Correlation between O*NET-derived theoretical exposure and Anthropic Economic Index observed adoption by degree field: $\rho \approx -0.08$ (not significant). This comparison is presented as an empirically interesting divergence rather than as validation of either measure. Multiple interpretations are possible, including: substantive temporal gap between capability and adoption; measurement artifacts from different levels of occupational aggregation; vendor-specific bias in adoption data derived from a single AI platform; and uneven enterprise governance and regulatory constraints. See main text Section 4.4 for full discussion.

Selected fields (sorted by theoretical exposure):

Field	Theoretical	Observed (% Claude usage)	Classification
Engineering Technology	0.463	2.35%	Latent vulnerability
Computer Science	0.399	26.11%	Adoption-led
Protective Services	0.384	0.20%	Latent vulnerability
Business	0.380	4.13%	Latent vulnerability
Education	0.231	10.65%	Adoption exceeds theory
Health Professions	0.232	1.52%	Latent vulnerability
Arts/Design/Media	0.220	9.06%	Adoption exceeds theory
Math/Statistics	0.271	26.11%	Adoption-led

Appendix E: PSEO Earnings Context

E.1 Coverage

The Census Bureau's Post-Secondary Employment Outcomes (PSEO) program links postsecondary student records to Unemployment Insurance wage records, providing median earnings at 1, 5, and 10 years post-completion by institution, degree field (CIP code), degree level, and graduation cohort. PSEO data are available for 34 participating states, covering 952 institutions total, of which 332 overlap with the analytical universe (43 R1, 31 R2, 32 Doctorate-level, remainder Master's, Baccalaureate, and Special Focus). PSEO earnings data are more comprehensive than College Scorecard earnings in that they are not limited to Title IV aid recipients, but coverage is geographically non-random due to the voluntary state participation model.

E.2 Field-Level Correlations (N = 4,649)

Comparison	Spearman ρ	p-value
AI exposure vs. 1-year median earnings	+0.257	< 0.0001
AI exposure vs. 5-year median earnings	+0.247	< 0.0001
AI exposure vs. 10-year median earnings	+0.233	< 0.0001
AI exposure vs. earnings growth (1yr→5yr)	+0.100	< 0.0001

E.3 Earnings by Field

CIP	Field	AI Exposure	1yr Median	5yr Median	10yr Median	1→5yr Growth

15	Engineering Technology	0.463	$49,492	$65,322	$76,015	32.0%
11	Computer Science	0.399	$40,342	$55,621	$67,482	37.9%
52	Business	0.380	$35,573	$46,768	$54,668	31.5%
14	Engineering	0.259	$35,452	$77,210	$97,356	117.8%
51	Health Professions	0.232	$58,854	$68,096	$75,733	15.7%
13	Education	0.231	$27,582	$39,778	$48,424	44.2%
42	Psychology	0.223	$28,566	$41,199	$53,578	44.2%
50	Visual/Performing Arts	0.220	$28,713	$39,143	$49,456	36.3%
26	Biological Sciences	0.198	$30,749	$49,757	$66,779	61.8%
27	Mathematics	0.271	$32,743	$61,121	$73,271	86.7%

E.4 Limitations

PSEO data are backward-looking (pre-AI-adoption earnings) and cannot speak to whether the earnings patterns observed will persist as AI adoption broadens. Coverage is non-random due to state-based participation. The 105-institution match to the analytical universe at the institution level is too small and too geographically concentrated for definitive institution-level benchmarking. The field-level comparison confirms that the exposure scores identify fields with high cognitive task density and current labor market premium, but it cannot confirm or disconfirm whether those fields will experience earnings disruption in the future.

Appendix F: Assumptions and Alternative Measurement Approaches

F.1 Task Classification Assumptions

The classification of 41 work activities into AI-positive, AI-negative, and neutral is the foundational judgment. Contested classifications:

"Getting Information" (classified AI-positive): Could be neutral if the operative skill is knowing *what* to seek rather than retrieval itself. Reclassifying would lower exposure for research-intensive occupations.

"Thinking Creatively" (classified neutral): Generative AI has destabilized this. Reclassifying as AI-positive would increase exposure for arts, design, and media — fields already showing high observed adoption (9.1% of Claude usage).

"Training and Teaching Others" (classified AI-negative): The Anthropic Economic Index shows education as the second-highest adoption category (10.7%). AI is being used *by* educators, not replacing teaching. The classification captures structural tasks but misses the augmentation dynamic.

F.2 Entry-Level Weighting Alternatives

No entry-level weighting (raw exposure): Produces 93.2% quadrant agreement with baseline — most stable alternative. Entry-level weighting refines but does not fundamentally reshape results.

Binary cutoff at Zone 3: Excludes Zone 4 occupations (accountants, financial analysts, software developers) entirely. Judged too aggressive.

Linear decay (Zone 1=1.0 to Zone 5=0.2): Produces >95% quadrant agreement. Step function retained for interpretive clarity.

F.3 CIP-SOC Crosswalk Alternatives

Census PSEO: Provides actual employment destinations for 332 institutions. Future version could replace crosswalk estimates for these schools.

Lightcast job posting data (proprietary): Would provide demand-side signal — are employers posting fewer openings in exposed occupations? Expected to show sharper geographic variation than WICHE state-level measures.

LinkedIn career pathways (not publicly available): Actual career transitions by degree field and institution. Would likely reveal that the crosswalk overstates pathway diversity.

F.4 Equal Weighting Rationale

Equal weighting gives AI exposure and demographic trajectory the same influence as historical backward-looking indicators. Revenue diversification (measuring tuition dependence — the ratio of net tuition revenue to total core revenues) is elevated; selectivity is retained with full weight. This is deliberate: the factor structure reflects historical relationships; equal weighting bets that emerging pressures will matter in the next decade.

The choice between equal and data-derived weights is the single most consequential methodological decision (61.3% quadrant agreement).

F.5 AI Exposure as Vulnerability vs. Opportunity

An ideal complement would measure institutional AI *adoption* in pedagogy — whether programs incorporate AI tools, update curricula, teach AI-augmented professional practice. No public dataset captures this. Proxies might include course catalog text analysis, institutional AI policy documents, or faculty AI use surveys.

F.6 Alternative Resilience Measures Considered

State appropriation trajectory. For public institutions, the multi-year trend in state appropriations per FTE student (available via SHEEO's State Higher Education Finance reports) is a direct resilience signal. Declining state support forces institutions to increase net tuition revenue to compensate, which affects affordability and enrollment — a cascading vulnerability. This data is available but adds a public-institution-specific measure with no private-institution equivalent, complicating cross-sector comparability.

Deferred maintenance burden. Institutions with large physical plants and accumulated deferred maintenance face capital expenditure demands that compete with operating budgets. This is a significant resilience factor for mid-tier regional institutions but is not consistently reported in IPEDS.

Debt service burden. Institutional long-term debt relative to total revenue affects financial flexibility and the ability to respond to enrollment shocks. West Virginia University's substantial construction-related debt was a precipitating factor in its 2023 financial exigency declaration and subsequent program eliminations. IPEDS reports some debt information but comparability across GASB and FASB reporting standards is limited.

Faculty composition. The ratio of tenure-line to contingent and adjunct faculty, and the age distribution of tenured faculty, affect both institutional cost structure and adaptive capacity. Institutions with large tenured faculties concentrated in declining fields face structural rigidity that constrains program reallocation. IPEDS reports faculty data through its Human Resources survey but integrating it adds complexity without clear analytical return at the institutional level.

F.7 Alternative Labor Market Measures Considered

Employer demand by field and region (Lightcast). Lightcast (formerly Burning Glass Technologies) maintains a comprehensive database of online job postings analyzable by SOC code, employer, required credentials, and metropolitan statistical area. Matching posting volume and trends to CIP-linked occupations would provide a forward-looking demand signal complementing the Scorecard's backward-looking post-college earnings data. This data requires an institutional subscription and was not available for this analysis.

Graduate school enrollment adjustment. For fields where the bachelor's degree is primarily a stepping stone to graduate or professional education (biological sciences, psychology, social sciences), median earnings at 1 or even 5 years post-completion understate the degree's long-run economic return because a substantial proportion of graduates are enrolled in further education rather than the labor market. A measure accounting for graduate enrollment rates by field would reduce the apparent labor market disadvantage of these pipeline fields.

Underemployment rate by field. The Federal Reserve Bank of New York publishes underemployment rates for recent college graduates by major field — the share working in positions that do not typically require a bachelor's degree. This is a more direct measure of post-college market position than earnings alone and would particularly affect the positioning of arts, humanities, and social science programs. Available at the national level only, not institution-specific.

Appendix G: Bifurcation Analysis Details

G.1 Distributional Statistics

- Resilience distribution skewness: 0.302
- Kurtosis: −0.097 (normal = 0, bimodal < 0)
- Bimodality coefficient: 0.376 (threshold for bimodality: 0.555)
- Middle tercile share: 33.4% (uniform would be 33.3%)

G.2 Mean Combined Score by Carnegie 2025 Tier

Tier	Mean	Median	N
R1	1.304	1.325	177
R2	1.070	1.073	121
RCU	1.058	1.000	180
Special Focus	1.028	1.021	275
Doctorate	0.965	0.982	187
Assoc/Bacc	0.916	0.909	56
Master's L/M	0.898	0.884	113
Master's S	0.888	0.894	143
Baccalaureate	0.878	0.873	303

G.3 Enrollment × Resilience Cross-Tabulation

Category	N	%
Growing + Above-median resilience	584	37.5%
Growing + Below-median resilience	196	12.6%
Shrinking + Above-median resilience	194	12.5%
Shrinking + Below-median resilience	581	37.3%

Supplementary Materials. v1.4. April 2026.