

THE FIRST RUNG MAP

Methodology and Data Documentation, Version 1.0

From the Universal HBCU Workforce Alignment Model

Sue Mukherjee, Ph.D. | Mukherjee & Associates LLC | Data vintage July 2026 | DOI: pending
Zenodo deposit

1. Purpose and scope

The First Rung Map is a free public instrument that shows, for a selected academic program family and state, three things no public instrument, to our knowledge and after a documented competitive review conducted in July 2026, presents together: the empirical occupational destinations of graduates in that field, the program family's earnings position against the state benchmark that anchors the federal earnings accountability framework effective July 1, 2026 (91 FR 40136), and the exposure of those destination occupations to artificial intelligence, including the distinction between occupations where AI augments work and occupations where it automates the entry-level tasks through which new graduates historically learned their jobs.

The instrument is the public expression of Systems 2 and 3 of the Universal HBCU Workforce Alignment Model, field-tested at Cheyney University of Pennsylvania and offered to the sector as a public good. It serves presidents, provosts, institutional research offices, deans, advisors, and students at the nation's HBCUs, and at any mission-driven institution navigating the same framework.

Version 1.0 covers the 15 to 20 program families, at the CIP family level, that account for the majority of HBCU bachelor's completions. Coverage expands in subsequent vintages.

2. What this instrument is, and is not

The First Rung Map presents national and state field-of-degree employment patterns. It does not present institution-specific graduate placements. No public dataset links individual institutions' completers to their occupations; the wage-record linkages that make such maps possible exist only inside state administrative data systems. Where an institution requires its own completer-level map, that analysis is a custom engagement, not a limitation this instrument hides.

The AI exposure values shown are measures of technical capability and observed usage, not predictions of job loss. Every source dataset carries this caveat, and this instrument carries it forward on every screen where exposure appears.

The earnings benchmark displayed is the published state threshold from the most recent federal calculation cycle, serving as a proxy until the Department of Education publishes benchmark values specific to the new framework. The two constructs differ, and Section 6 states exactly how.

3. Data sources and vintages

Layer	Source	Vintage	Access and license
Occupational fans	ACS Public Use Microdata Sample, 5-year file, via IPUMS USA	Most recent 5-year file	Public; IPUMS citation required
Program-to-field mapping	Instrument-published FOD-to-CIP-family mapping (Section 4)	Version 1.0	Published methodology artifact of this instrument
Earnings benchmark	Federal earnings thresholds by state, 89 FR 107130 (CY2024), derived from 2022 ACS 5-year PUMS	December 31, 2024	Public domain (Federal Register)
Program earnings margin	College Scorecard field-of-study files, median earnings 4 years after completion, 4-digit CIP by credential level	Updated June 10, 2026	Public domain (U.S. Department of Education)
Program weights	IPEDS Completions survey, HBCU institutions by CIP	Most recent final release	Public domain (NCES)
AI exposure	GPTs are GPTs occupation-level exposure scores (Eloundou, Manning, Mishkin, Rock), occ_level.csv	2023, revised 2024	MIT license
AI applicability	Working with AI applicability scores (Tomlinson, Jaffe, Wang, Counts, Suri), Microsoft Research	2025	CC BY 4.0
Automation vs augmentation	Anthropic Economic Index public data releases	Most recent release	CC BY with attribution
Entry-level employment signal	Brynjolfsson, Chandar, and Chen, Canaries in the Coal Mine, Stanford Digital Economy Lab	Revised November 13, 2025	Cited finding, not embedded data
Mobility context	Opportunity Insights Mobility Report Cards	Cohorts 1999 to 2013	Public for research with attribution
Certificate alignment	State Workforce Pell occupation lists under P.L. 119-21	State by state, 2026	Public (state agencies)

4. How the map produces a result

The pipeline runs: program family, to the published FOD-to-CIP-family mapping, to the ACS occupational fan, to the earnings distribution, to the AI exposure and applicability overlay, to the first-rung heuristic, to the displayed result. Each output belongs to one of three epistemic categories, and the interface labels each accordingly.

Computed outputs: the occupational fan, the earnings margin, and program completion weights. These are calculated from the named public datasets under the stated suppression rules.

Matched external indicators: AI task exposure, AI applicability, automation-versus-augmentation lean, and mobility context. These are values published by external researchers, joined to this instrument's occupations and institutions without modification.

Interpretive flags: the first-rung flag, the judgment intensity index, and mission value framing. These are classifications this instrument constructs from the evidence above. They are labeled as interpretation wherever they appear, and their construction rules are stated in this document.

5. The occupational fan: construction and honest limits

Each fan is built from the ACS 5-year PUMS by cross-tabulating field of bachelor's degree (DEGFIELD, DEGFIELD2) with occupation (OCC) and earnings (PINCP, WAGP), weighted by person weight (PWGTP), with variance estimated by the 80 successive-difference replicate weights. Early-career views use ages 22 to 27, following the Federal Reserve Bank of New York's convention for its recent-graduate series; prime-age views use ages 25 to 54.

No official crosswalk between the Census field-of-degree taxonomy and the NCES Classification of Instructional Programs exists. The Census Bureau's own documentation states the classifications are related but not identical. This instrument therefore publishes its own FOD-to-CIP-family mapping as a versioned, downloadable methodology artifact, and operates at the CIP family level rather than claiming 6-digit precision it cannot support. When an official crosswalk is published, this instrument adopts it.

Suppression rule: any fan cell with a coefficient of variation above 30 percent, or an unweighted sample below 50 persons, is suppressed or flagged. Views restricted to Black bachelor's holders apply the same rule and, at state level, frequently require CIP-family pooling; the instrument states when a displayed value reflects pooled data.

What the fan cannot show: institution-specific placements, completion cohort timing, or Title IV receipt. The fan describes where the people who hold a degree in this field actually work and what those roads pay. That is the honest construct, and it is labeled as such on every view.

6. The earnings margin: two floors, never conflated

The federal framework effective July 1, 2026 compares an undergraduate program's completer median earnings, measured four years after completion, against the median earnings of working adults aged 25 to 34 whose highest credential is a high school diploma, in the institution's state. The Department of Education has not yet published the state benchmark values specific to this framework.

The instrument therefore displays the published Calculation Year 2024 thresholds from the prior gainful employment framework (89 FR 107130) as the current best proxy, labeled exactly that way. The two floors differ in universe: the prior framework's threshold includes unemployed labor force participants, while the new framework's universe is working adults not enrolled during the measured

year. The new floor will generally sit somewhat higher. Every margin shown carries the label identifying which floor it uses, and the instrument swaps in the official values the day the Department publishes them.

Program-side earnings come from College Scorecard field-of-study files at the 4-digit CIP level. Cells suppressed by the Department for small cohorts display as suppressed, never estimated. The margin shown is an estimate for orientation and planning; it is not the official determination, which only the Department's calculation produces.

7. The AI layer: exposure, applicability, and the first rung

The source datasets use related but distinct occupational taxonomies: the ACS codes occupations in OCC, the federal standard is SOC 2018, and the exposure datasets publish at O*NET-SOC or SOC levels. The instrument harmonizes through the published Census OCC-to-SOC crosswalk and the O*NET-SOC-to-SOC correspondence, at the most detailed level a match supports. Where an occupation cannot be matched across taxonomies, the instrument displays not covered for the unmatched layer rather than borrowing a neighboring occupation's value.

For each destination occupation in a fan, the instrument displays three values with primary citations: task exposure (Eloundou et al.), AI applicability (Microsoft Research), and automation-versus-augmentation lean (Anthropic Economic Index). Where an occupation is absent from a source dataset, the instrument displays not covered rather than an imputed value.

The layer separates two questions public discussion routinely merges: whether the occupation is healthy, and whether its entry rung remains reachable. The Stanford Digital Economy Lab finds early-career workers aged 22 to 25 in the most AI-exposed occupations experienced a relative employment decline of roughly 13 to 16 percent depending on specification and data vintage, concentrated in occupations where AI automates rather than augments, with adjustment through employment rather than pay. As supporting labor-market interpretation rather than primary causal evidence, PwC's 2026 Global AI Jobs Barometer finds AI-exposed entry-level postings are seven times more likely to require traditionally senior skills, with such seniorised entry postings growing 35 percent since 2019 while other entry roles shrank 10 percent. In the same supporting category, the Burning Glass Institute's Expertise Upheaval framework distinguishes roles where AI compresses the learning curve from roles where expertise must still be accrued. The instrument's first-rung flag is a clearly labeled interpretive classification, not a computed prediction. The flag is assigned when an occupation combines three conditions: high AI task exposure in the matched external indicators, high routine cognitive task content in its O*NET activity profile, and published evidence of early-career pressure or seniorised entry postings in that occupation or its group. The inputs are named here even though Version 1.0 publishes no numeric weights; readers can inspect every input the flag draws on.

Correlation is not causation. The Stanford dashboard's own language is the model: these measures describe the correlation between occupational exposure and employment trends, not a causal link. This instrument adopts the same posture.

8. Mission value evidence

The Hope Map view carries institutional mobility context from the Opportunity Insights Mobility Report Cards, which reflect students who attended college between 1999 and 2013. These are structural evidence of mission, not current outcome measures, and are labeled accordingly. Sector context draws on UNCF's finding that HBCUs move students from the bottom 40 percent of household income into the middle class at a 34.3 percent rate against 15.8 percent nationally. The framework's orienting concept, serviceability over employability, is Dr. Lezli Baskerville's, and is credited to her wherever it appears.

9. Judgment intensity

The judgment dimension follows the published construction of NBER Working Paper 28733 (Deming, April 2021), verified against the paper directly: a simple, unweighted average of three O*NET work activities, Making Decisions and Solving Problems, Developing Objectives and Strategies, and Planning and Prioritizing Work, rescaled to a 0 to 10 scale. The paper does not specify importance versus level ratings, so this instrument states only that O*NET work activity ratings are used. Equal weighting is retained because it is the published precedent; this instrument invents no weights of its own. The index is presented as a proxy for the human judgment content of an occupation, one signal among several, not a validated measure of AI resistance.

10. Versioning, citation, and error reporting

The instrument, its dataset, and this document are versioned together. Each data vintage receives a dated changelog entry and a new version DOI under a persistent concept DOI. Suggested citation: Mukherjee, S. (2026). The First Rung Map, from the Universal HBCU Workforce Alignment Model [data instrument], Version 1.0. Mukherjee & Associates LLC. suemukherjee.com/first-rung-map. Errors may be reported to sue@suemukherjee.com and are acknowledged in the changelog.

11. Accessibility

The instrument targets WCAG 2.1 AA conformance: full keyboard navigation, labeled numeric alternatives for all controls, results never encoded by color alone, and text summaries accompanying every chart.

12. Limitations, stated plainly

Field-of-degree patterns are not placements. The benchmark shown is a labeled proxy pending official values. AI exposure measures capability and usage, not displacement. Mobility data are historical. The FOD-to-CIP mapping is approximate and published for inspection. Small cells are suppressed, not smoothed. Where this instrument cannot support a claim, it does not make one.

Acknowledgments: Advisory reviewers to be named upon their confirmation. The serviceability framing is Dr. Lezli Baskerville's. The Universal HBCU Workforce Alignment Model was field-tested at Cheyney University of Pennsylvania, the nation's first HBCU, under the leadership of President Aaron A. Walton.