Original Research Article

AI Task Exposure and U.S. Employment Patterns Across Demographic, Educational and Occupational Groups, 2022–2025

Shanmukh Sriram. Chintapalli[1] and Dr. Madhura Shivaram[2]

[1]Sri Kumarans' Children's Home, Survey No 44 – 50, Mallasandra Village, Uttarahalli Hobli, Bengaluru South Taluk, Karnataka 560109, India; [2]Independent Researcher, No.779 A block, 14th Main, Sahakara Nagara, Bengaluru , Karnataka, 560092, India

ABSTRACT

Since ChatGPT launched in November 2022, artificial intelligence (AI) has rapidly entered the workplace, raising questions about whether its effects are felt equally across different groups of workers. This paper analyzes observed 2022–2025 employment associations alongside BLS 2024–2034 occupational projections to characterize both current patterns and anticipated trends, associations between occupational AI task exposure and employment patterns across demographic groups – and which groups face the greatest risk based on projected trends. Using U.S. government employment survey data from 2022 to 2025, jobs are classified into three categories: those with higher measured AI task exposure where substitution risk may be elevated, those where AI tools appear associated with productivity gains, and those showing minimal measured AI task exposure. A machine learning model identifies which factors best predict whether a job category will grow or shrink. Overall employment remained stable, but this masks meaningful differences: women make up around 70% of workers in roles with the highest AI task exposure scores, concentrated in office and administrative work, where substitution risk may be elevated; workers without a college degree show nearly zero projected job growth, while those with advanced degrees are projected to gain over 9%. Education level and occupation type are closely associated with changes in employment. This paper concludes that employment patterns associated with AI exposure do not appear equal across all workers – observed patterns differ by occupation and education. Because the study covers a short time window, some trends may reflect post-pandemic recovery rather than AI alone, and future research should apply stronger causal methods over longer timeframes.

Keywords: Artificial intelligence; labor market; employment; AI exposure; gender; education level; occupation; machine learning

INTRODUCTION

From 1995 to 2022, US labor markets endured five distinct cycles – the dotcom bust, the 2008 Global Financial Crisis, a gradual recovery through the 2010s, the COVID-19 pandemic collapse, and a post-pandemic rebound – each reshaping employment levels, participation rates, and workforce composition (1–3). By 2022, headline unemployment had returned to pre-pandemic levels near 3.5–4.0%, though labor

Corresponding author: Shanmukh Sriram. Chintapalli, E-mail: shanmukh.chintapalli@gmail.com.

Accepted August 11, 2026
https://doi.org/10.70251/HYJR2348.449941015

force participation remained below its pre-2008 trend, setting the context for the 2022–2025 window this paper examines.

Since ChatGPT's public release in November 2022, aggregate US labor market indicators have not shown a measurable decline that can be directly linked to AI adoption (4). Some economists note that employment statistics typically lag structural shifts by several years, making near-term macro data an incomplete basis for assessing displacement (5). On the ground, however, the signals are more pointed: nearly 55,000 job cuts in 2025 were cited by employers as associated with AI adoption, with companies including Amazon, Microsoft, Workday, and Salesforce explicitly citing AI when announcing large-scale workforce reductions (6, 7). Research suggests that rather than solely eliminating jobs, employment growth in augmentation-prone roles coincides with increased AI adoption; clerical, data-entry, and routine cognitive roles face a declining employment share (5, 8, 9).

Despite widespread AI adoption, there is limited empirical evidence on whether AI-related employment changes affect all workers equally. This is the central question: is AI Task Exposure differentially associated with employment outcomes across demographic groups – and which groups appear most exposed to risk, with education as a key moderating variable across every AI Exposure tier. This paper addresses that gap by analysing demographic and occupational differences in employment trends associated with AI task exposure using U.S. employment data from 2022-2025. To examine the shift, this paper draws on the Current Population Survey (CPS) via IPUMS and Bureau of Labor Statistics (BLS) data, combined with Anthropic's AI Exposure Index cross-walked against CPS occupational codes, to capture differential vulnerability and opportunity across the workforce (3, 10–12). The analysis spans gender, age, race, educational attainment, and occupation, with AI Exposure layered across each dimension, a Gradient Boosting Model framework modelling associations between employment change between 2022 and 2025 as a function of AI Exposure and demographic and education level attributes. The GBM (Gradient Boosting Machine) classifier (Area Under the Curve, AUC 0.56) demonstrates stronger predictive signal for employment decline than growth; SHAP (Shapely Additive Explanations) and interaction analyses provide directional evidence that occupation-level AI adoption is associated with differential labor market exposure across demographic groups.

Findings reveal that AI's association with the labor market appears uneven and still evolving; observed shifts appear disproportionately concentrated among demographics represented in clerical, entry-level, and lower-credentialed roles. Drawing on prior work (13), this study disaggregates employment levels across gender, race, age, and education simultaneously, mapping those dimensions onto repeated cross-sectional CPS aggregates (2022–2025) to assess where AI Substitution Risk is concentrating (8, 13).

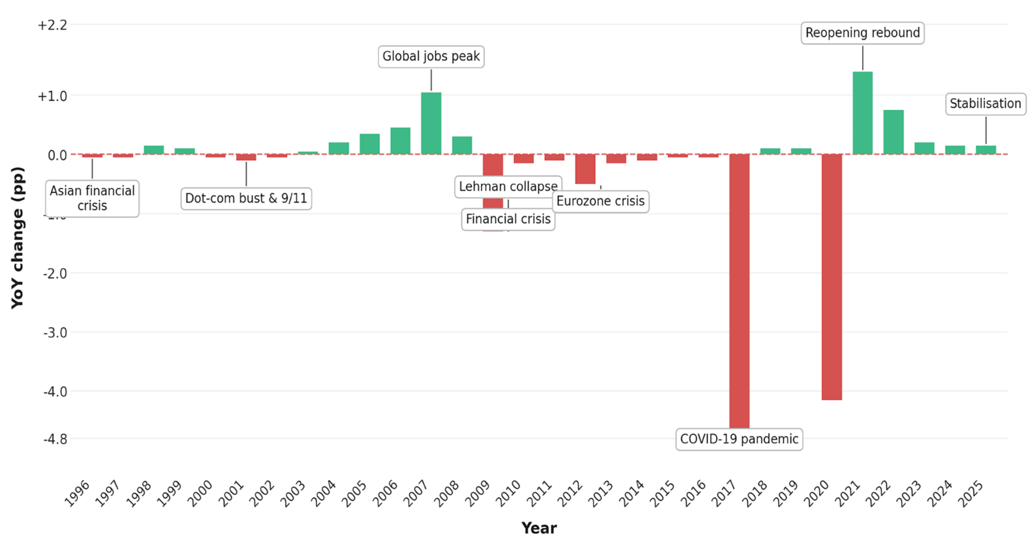

Figure 1. *U.S. Employment-to-Population Ratio (%) and unemployment rate (%), 1996–2025, from published BLS CPS time series; contextualizes the 2022–2025 study window against long-run labor-market history.*

METHODS AND MATERIALS

Data Collection For The Post-ChatGPT Era Analysis

IPUMS CPS is the standard for U.S. labor force analysis – it is the official microdata source underlying BLS headline statistics, ensuring that weighted aggregations are nationally representative and institutionally credible (3, 11, 12). Its harmonised SOC-coded occupational coverage across gender, race, age, and education in a single harmonized dataset enables the three-way crosswalk with Anthropic's AI Exposure Index that is the analytical core of this paper (10). This publicly available survey delivers the combination of demographic depth, occupational granularity, and consistency at scale, enabling the detailed analysis which is central to this research effort.

Primary Data Source

This analysis draws on microdata from the Integrated Public Use Microdata Series Current Population Survey (IPUMS CPS), administered by the Minnesota Population Centre at the University of Minnesota (3). The Current Population Survey (CPS) is a monthly household survey jointly sponsored by the U.S. Census Bureau and the Bureau of Labor Statistics (BLS), providing the principal source of labor force statistics for the civilian noninstitutional population of the United States (14). IPUMS harmonises CPS microdata across survey years, applying consistent variable codes and documentation to facilitate longitudinal analysis (3). Data were extracted for survey years 2022 onward–a period coinciding with the accelerating mainstream diffusion of generative artificial intelligence technologies–enabling examination of how AI Exposure patterns relate to contemporaneous employment outcomes.

Variables Extracted

A custom data request was submitted through the IPUMS CPS online portal. Eleven variables were downloaded spanning survey-year identification, two population weights, four demographic characteristics, two labor market status indicators, an occupation code, and an educational attainment recode. Table 1 summarises each variable, its data type, value range, and analytical role (3).

All weighted aggregations use COMPWT as the weighting field, as it is calibrated to replicate official BLS labor force totals and accounts for the rotating panel

Table 1. *CPS variables extracted from IPUMS CPS, 2022–2025: survey year, two weights (WTFINL, COMPWT), four demographics (age, sex, race, Hispanic origin), labor-force status, occupation code, and education recode.*

Variable	Label	Category / Values	# Distinct Values	Description
YEAR	Survey year	Integer (2022–2025)	4	Temporal anchor; 2022 onward used
WTFINL	Final Basic Weight	Continuous (≥ 0)	–	Poststratified inverse-probability weight
COMPWT	Composite Weight	Continuous (≥ 0)	–	Primary weight; replicates BLS labor force totals
AGE	Age	Integer (0–85+)	86	Recoded into 15 standardised age buckets
SEX	Sex	Male; Female	2	Binary grouping dimension
RACE	Race	White (100); Black (200); American Indian/Aleut/Eskimo (300); Asian only (651); Hawaiian/Pacific Islander only (652); WhiteBlack (801); WhiteAmerican Indian (802); WhiteAsian (803); WhiteHawaiian/Pacific Islander (804); BlackAmerican Indian (805); BlackAsian (806); BlackHawaiian/Pacific Islander (807); American IndianAsian (808); etc (809–830)	25	Detailed IPUMS race classification (raw extraction level); combined with HISPAN into mutually exclusive race/ethnicity groups

Continued Table 1. *CPS variables extracted from IPUMS CPS, 2022–2025: survey year, two weights (WTFINL, COMPWT), four demographics (age, sex, race, Hispanic origin), labor-force status, occupation code, and education recode.*

Variable	Label	Category / Values	# Distinct Values	Description
HISPAN	Hispanic origin	Not Hispanic (000); Mexican (100); Puerto Rican (200); Cuban (300); Dominican (401); Salvadoran (411); Central American excluding Salvadoran (412–419); South American (420–429); Other Hispanic (500)	9	Detailed IPUMS Hispanic origin; combined with RACE into mutually exclusive race/ethnicity groups
EMPSTAT	Employment status	At work; Has job/not at work; Unemployed; NILF	5	Basis for #Employed and #Unemployed counts in final data set
LABFORCE	Labor force status	In labor force; Not in labor force; N/A	3	Basis for #Not_In_Labor_Force count in final data set
OCC	Occupation code	Integer SOC codes (0–9840)	~500	Linked to AI Exposure Score via SOC
EDUC	Educational attainment	No schooling → Doctoral degree	14	IPUMS harmonised recode variable

structure of the CPS sample design (11, 14). WTFINL was retained for reference and sensitivity checks only.

Data Processing and Age Recoding

A Python script aggregated individual CPS records into population-weighted cells defined by all categorical dimensions jointly: YEAR, SEX, RACE, HISPAN, OCC, EDUC, and AGE bucket. Within each cell, COMPWT values were summed conditional on employment status to produce three mutually exclusive outcome counts: #Employed (EMPSTAT 'at work' or 'has job not at work'), #Unemployed (EMPSTAT 'unemployed'), and #Not_In_Labor_Force (LABFORCE 'not in labor force') (14). The continuous AGE variable was recoded into fifteen five-year brackets following BLS tabulation conventions: under 16; 16–17; 18–19; 20–24; 25–29; 30–34; 35–39; 40–44; 45–49; 50–54; 55–59; 60–64; 65–69; 70–74; and 75 and over. Cells in which all three outcome counts equal zero are retained for completeness but excluded from rate calculations.

Detailed IPUMS RACE (25 categories, including all two, three, four, and fiverace combinations such as 'American IndianAsian' and 'WhiteBlack') and HISPAN (9 categories of detailed Hispanic origin) were combined into mutually exclusive race/ethnicity groups as follows: respondents with any Hispanicorigin HISPAN code were classified as 'Hispanic or Latino' regardless of RACE; among nonHispanic respondents, singlerace RACE codes were mapped onetoone to White, Black or African American, Asian, Native Hawaiian or Other Pacific Islander (NHPI), or American Indian or Alaska Native (AIAN); nonHispanic respondents reporting two or more races (including the 'American IndianAsian' code) were classified as Multiracial.

AI Exposure Index and AI Exposure

Occupation-level AI Exposure scores were sourced from Anthropic's economic impact report, which uses a task-based methodology to evaluate the degree to which each SOC-coded occupation is exposed to automation or augmentation by large language model AI systems (10). Each occupation's tasks–drawn from the O*NET database–are evaluated for AI accessibility, and the results are aggregated to a continuous AI Exposure Score bounded between 0 and 1 (10, 15). A score of 0 indicates no meaningful AI Exposure (tasks require physical manipulation or real-world sensory engagement beyond current LLM capabilities), while a score of 1 indicates maximal exposure (the vast majority of task content is accessible to current AI systems). Intermediate scores reflect partial exposure: scores near 0.3 indicate a minority of tasks are AI-accessible; scores near 0.5 indicate approximately half; and scores near 0.75 indicate a substantial majority. Knowledge-intensive roles such as legal analysts, financial advisors, and software developers typically score closer to 1, whereas personal care workers, construction trades, and emergency responders typically score at or near 0 (10). For the purposes of this analysis, each occupation carries two independently-sourced fields from Anthropic's occupational AI Exposure dataset: a continuous AI_Exposure_Score (0–1) reflecting the overall magnitude of task-level AI accessibility, and a categorical AI_Exposure_Label – NO_AI_EXPOSURE, AI_COMPLEMENT, or AI_DISPLACEMENT – generated by prompting Antropic Claude to classify each occupation's dominant usage pattern. The label is not a fixed numeric recoding of the score; it reflects the predominant character of how AI is used within an occupation's task mix (automation, augmentation, or minimal current engagement) as assessed by the model, which need not scale linearly with the aggregate exposure magnitude (10). Both fields are reported in Table 2 and used throughout the Results as complementary measures – the score for magnitude, the label for usage pattern.

The NO_AI_EXPOSURE designation identifies occupations where tasks require physical manipulation or real-world sensory engagement largely outside current LLM capabilities – examples include personal care workers, construction trades, and emergency responders.

The AI_COMPLEMENT designation identifies occupations in which AI adoption is associated with task augmentation, where AI handles discrete tasks while core professional judgement remains with workers – examples include Education and Childcare Administrators and Medical and Health Services Managers.

The AI_DISPLACEMENT designation identifies occupations in which a substantial share of task content is measurably accessible to current AI systems, corresponding to elevated potential substitution risk; this does not imply that job loss will occur – examples include Office and Administrative Support and Sales and Related occupations.

Final Microdata Structure

The CPS aggregate data was then joined with the AI Exposure scores data to create a crosswalk between the two datasets. After merging CPS aggregates with AI Exposure Scores via the OCC field, the final dataset contains thirteen columns: YEAR, SEX, RACE, HISPAN, OCC,OCC_DESCRIPTION, EDUC, AGE_BUCKET, AI_Exposure_Label, AI_Exposure_Score, #Employed, #Unemployed, and #Not_In_Labor_Force. Table 2 presents four sample observations from survey year 2022 immediately following the join of the two tables and further aggregations have been done during trend analysis.

Data Limitations

Claude AI usage: The AI Exposure index is based primarily on Claude usage data and does not include usage of other AI tools such as ChatGPT or Gemini.

Cross-walk coverage: The cross-walk resulted in ~95% coverage of AI Exposure index vs. the Occupation codes from BLS data.

Automation vs. augmentation ambiguity: The index estimates the share of queries that are automation

Table 2. *Four sample 2022 records from the merged dataset: IPUMS CPS demographics joined with Anthropic AI Exposure data by occupation code, showing weighted employed, unemployed, and not-in-labor-force counts. AI_Exposure_Score (continuous magnitude) and AI_Exposure_Label (Claude-generated usage-pattern category) are independently sourced fields.*

YEAR	SEX	RACE	HISPAN	OCC	OCC_DESCRIPTION	EDUC
2022	Female	Asian only	Not Hispanic	1021	Software developers	Master's degree
2022	Female	Asian only	Not Hispanic	4230	Maids and housekeeping cleaners	High school diploma or equivalent
2022	Female	Asian only	Not Hispanic	4720	Cashiers	High school diploma or equivalent
2022	Female	Asian only	Not Hispanic	4720	Cashiers	Some college but no degree

AGE_BUCKET	AI_Exposure_Label	AI_Exposure_Score	#Employed	#Unemployed	#Not_In_Labour_Force
35 to 39 years	AI_COMPLEMENT	0.288	236764	3950	3600
30 to 34 years	NO_AI_EXPOSURE	0	21732	319	431
25 to 29 years	AI_DISPLACEMENT	0.085	27150	13770	983
20 to 24 years	AI_DISPLACEMENT	0.085	71310	2199	7482

related, augmentative, or neither, used as a proxy for whether AI usage is complementary or substitutable with labor – but these categories are model-assigned and not directly observed.

The term 'AI Displacement,' as used in this paper, describes occupations whose core tasks are largely executable by current AI tools. It measures the degree of AI Task Exposure within an occupation, and should not be interpreted as a prediction of actual job loss.

Ethics Statement

This study relies exclusively on secondary, publicly available, de-identified microdata (IPUMS CPS, U.S. Bureau of Labor Statistics tabulations, and O*NET occupational data) and does not involve direct contact with, or collection of new data from, human participants. No personally identifiable information was accessed or processed at any stage of the analysis. As the research involved no human-subjects contact and used only de-identified, publicly released government data, it is considered exempt from Institutional Review Board (IRB) review; no formal IRB submission was made for this project.

RESULTS

This section examines employment and unemployment trends across key demographic dimensions. Data is analysed by gender, race/ethnicity, and age group to identify patterns and disparities in labor market outcomes. These breakdowns provide a comprehensive view of workforce participation, highlighting how economic conditions and opportunities differ across population segments over time.

Trends by Gender

Table 3 shows the general employment trend. The Employment-to-Population Ratio (EPR) for males consistently exceeded that of females by approximately 10 percentage points (pp) – 65.6% for men versus 54.8% for women in 2022, narrowing slightly to 65.0% and 55.0% respectively in 2025.

Overall, compared to 2022, the employment percentage of both males and females with an advanced degree has increased. While there is a 14 percentage point gap between male and female employment rates in the less than high school degree category, this gap narrows to 3 percentage points at the advanced degree level. AI Exposure classifications – AI Displacement, AI Complement, and No AI Exposure – are mapped from CPS three-digit OCC codes using published AI task-content frameworks (10, 15). Female workers are disproportionately represented in several occupations with high AI task Exposure scores, which are associated with elevated substitution risk.

Key Findings

1. The male–female employment gap is persistent but the absolute employment gap is widening. Table 3 shows that male employment rose from 81.5M to 86.6M between 2022 and 2025, while female employment grew from 73.5M to 76.9M over the same period – a male gain of 5.1M versus a female gain of 3.4M. The table also shows that the Employment-to-Population Ratio (EPR) gap, which stood at 10.8 pp in 2022 (65.6% male vs. 54.8% female), narrowed only marginally to 10.0 pp by 2025. Both the rate gap and the absolute count gap point to persistent differences in employment growth trajectories across genders. The two-proportion z-test gave $z=2.15$ ($p=0.032$), showing the gap narrowing is statistically real, not sampling noise.
2. Female employment is highly concentrated in a narrow band of occupations. Figure 2 reveals a stark split: six occupations register female shares of 69–83%, while eight occupations sit at 27% or below. Healthcare support (83%), personal

Table 3. *Employment-to-Population Ratio (%) and employed population (millions) by sex, 2022–2025, with year-over-year and male–female gap changes (pp). IPUMS CPS, weighted; gap narrowing z=2.15, p=0.032.*

Year	Male employed (M)	Male EPR %	YoY Δ	Female employed (M)	Female EPR %	YoY Δ	M–F gap
2022	81.5	65.6%	–	73.5	54.8%	–	10.8 pp
2023	83.4	65.7%	+0.1	74.3	55.4%	+0.6	10.3 pp
2024	85.3	65.3%	−0.4	76.0	55.4%	0.0	9.9 pp
2025	86.6	65.0%	-0.3	76.9	55.0%	−0.4	10.0 pp

care and service (77%), healthcare practitioners (75%), education (74%), office and administrative support (70%), and community and social services (69%) account for the bulk of female employment. By contrast, construction and extraction (4%), installation and repair (5%), architecture and engineering (19%), protective services (23%), and transportation (21%) remain heavily male-dominated. This concentration into a small cluster of high-female-share occupations shapes every other pattern observed across the section.

3. Office and administrative support is the single most exposed female-dominated occupation. Figure 3 plots AI Displacement level against female occupational share, with bubble size reflecting workforce scale. Office and administrative support – 70% female and approximately 10% of the total workforce – sits at the upper-right of the chart with an AI task Exposure score of 32.8%, the highest among large female-dominated groups. Business and financial operations and legal occupations also carry higher AI task Exposure scores (both above 20%) and sit near or above the average female share of 44.7%. By contrast, the largest female-dominated occupations – healthcare practitioners, healthcare support, and personal care – cluster near 0% on the AI task Exposure scale at the far right of the chart.

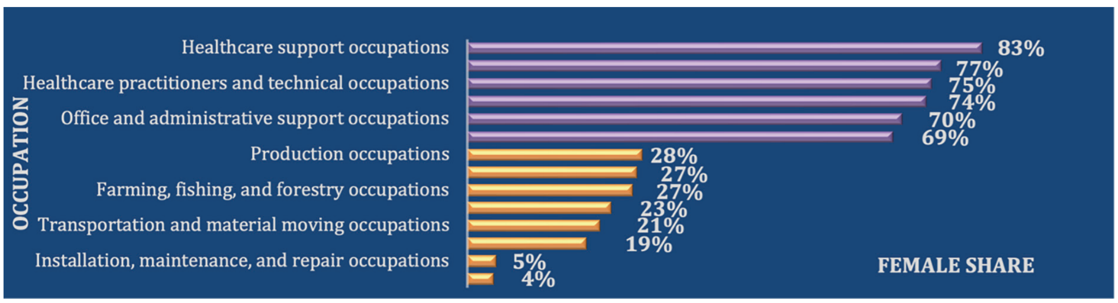

Figure 2. *Female employment share (%) by major occupation group, pooled 2022–2025. Share = weighted female workers ÷ total weighted workers per group. Source: IPUMS CPS, weighted by COMPWT.*

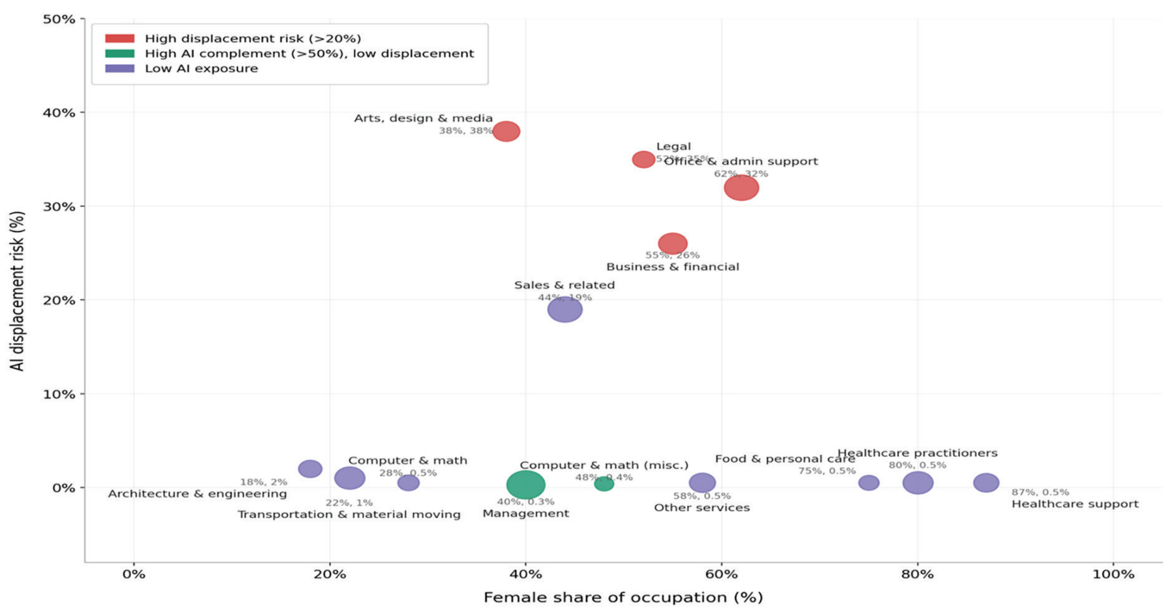

Figure 3. *AI Task Exposure score (%) vs. female employment share (%) by occupation group; bubble size = group's share of total workforce. Source: IPUMS CPS merged with Anthropic AI Exposure Index.*

4. Higher education lifts EPRs for both genders but narrows the gap most at the top credential tier. Table 4 shows that at the "less than high school" level the gender EPR gap was 15.6 pp in 2022 (45.9% male vs. 30.3% female), and Figure 4(B) confirms it narrowed only slightly to 14.0 pp by 2025. The gap compresses progressively up the credential ladder, reaching 3.0 pp at the advanced degree level (73% male vs. 70% female in 2025, per Figure 4(A)). However, The two-proportion z-test gave z=0.90 (p=0.37), indicating the gap narrowing at this tier is not significant. Figures 5(a) and 5(b) show that women with advanced degrees recorded the largest percentage-point gain of any group: +10.5 pp versus +11.5 pp for men at the same level. Notably, females with less than a high school degree and with some college saw small declines in EPR (−1.4% and −2.1% respectively), a divergence not observed for males at those same levels.
5. Women are disproportionately concentrated in occupations categorized as high AI Exposure roles. Figure 5 shows the share of employment by gender across the three AI Exposure tiers. In the AI Displacement category, women account for 70% of workers versus only 30% for men – a more than two-to-one skew. This contrasts sharply with the AI Complement category (46% female, 54% male) and the No AI Exposure category (44% female, 56% male), both of which show near-parity. The chi-square test gave χ^2=2799.9 ($p<0.001$), confirming sex and exposure tier are strongly associated, not chance.

Table 4. *Percentage change in employment, 2022–2025, by sex and educational-attainment tier (ages 25+), based on population-weighted employed counts. Source: IPUMS CPS.*

Education level (25+)	% Change in Employment 2022-2025	
	Male	**Female**
Less than High school	+0.8%	-1.4%
High School graduation	+6.2%	+2.8%
Some college	+1.4%	-2.1%
Associate / Professional Degree	+9.7%	+4.9%
Bachelor's degree	+7.1%	+8.2%
Advanced Degree	+11.5%	+10.5%
Overall	+6.2%	+4.7%

A

B

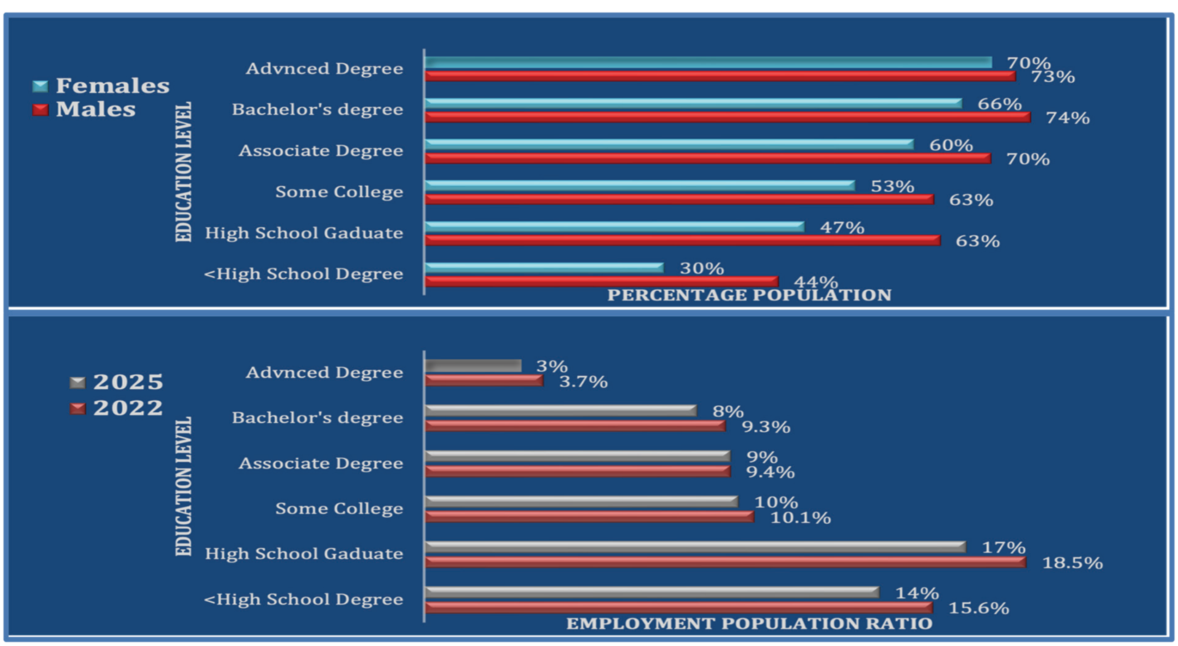

Figure 4. *Employment-to-Population Ratio (EPR) by Educational Attainment and Sex, IPUMS CPS, 2022–2025 (A)Top: EPR (%) by educational attainment and sex, pooled 2022–2025; gap narrows from 14 pp at "<High School" to 3 pp at "Advanced Degree." (B) Male–female EPR gap (percentage points) by educational attainment, comparing 2022 vs. 2025; gap narrowed slightly across most tiers over the period.*

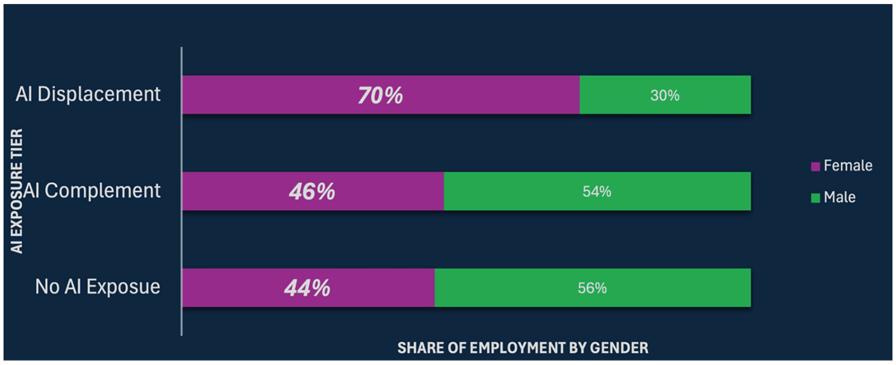

Figure 5. *Share of workers (%) by sex across three AI Exposure tiers – Displacement, Complement, No Exposure – 2022–2025. Source: IPUMS CPS + AI Exposure Index.*

Employment Trend by Race

Disaggregating employment outcomes by race and ethnicity is essential because these groups differ systematically in their occupational distributions, educational credential profiles, geographic concentration, and historical exposure to labor market discrimination – factors that interact with macroeconomic shocks and shifts, such as the diffusion of generative AI, in ways that group-level averages cannot capture. Without this disaggregation, potential associations between high AI exposure and employment shifts, if concentrated among specific communities, may remain invisible in headline Employment-to-Population Ratios, delaying policy recognition and response. The 2022–2025 window makes this analytical lens particularly urgent: as AI Exposure scores rise unevenly across occupations, the racial and ethnic groups most concentrated in high Exposure roles may experience materially different near-term labor market trajectories than those whose occupational mix has low measured exposure to current generative AI capabilities. Analysing employment trends across the race/ethnicity categories reported in Figure 6A – White (nonHispanic), Black or African American, Hispanic or Latino, Asian, and Other (a combined category comprising American Indian/Alaska Native, Native Hawaiian/Other Pacific Islander, and Multiracial respondents, pooled because of small unweighted cell counts; see Methods) – is therefore not merely a demographic exercise but a relevant exercise for understanding where the emerging labor-market adjustments are associated with AI-intensive occupational exposure Asian teens have the lowest Employment-to-Population Ratio (19%), while Asian workers aged 55 and above have the highest Employment-to-Population Ratio. Among population in Prime Working Age, Black males and Hispanic females have the least Employment-to-Population Ratio, 79% and 70% respectively. Hispanic or Latino workers have the highest share in construction, building and grounds cleaning, and farming & fisheries occupations. Black or African American workers have the highest share in healthcare support occupations. Asian workers have the highest share in computer and mathematics occupations.

<u>Key Findings:</u>

1. Hispanic EPR shows a modest but statistically non-significant narrowing relative to White and Asian workers: Table 6A shows that White (non-Hispanic) and Asian workers entered the period with the highest EPRs, hovering near 63%. Hispanic or Latino workers raise from 62.1% in 2022 to 63.0% in 2024. By contrast, Black workers lagged White workers by approximately 4 to 5 percentage points throughout the period, with no comparable trajectory of improvement visible in the data. However, the WLS trend test gave $t=0.89$ ($p=0.54$), showing Hispanic EPR is flat, not rising;
2. Asian workers exhibit the widest within-group age employment disparity of any racial group. Figures 6A together reveal a striking pattern unique to Asian workers: they record both the lowest teen Employment-to-Population Ratio of any racial group at 19%, and the highest Employment-to-Population Ratio among older workers aged 55 and above. No other racial or ethnic group spans both extremes simultaneously. Figure 6C shows that 25% of Asian workers are employed in computer and mathematics-related occupations. The two-proportion z-test gave $z=-18.79$ ($p<0.001$), confirming the age disparity is real.
3. Within the prime working-age population, the

gender–race intersection produces the sharpest employment gaps. Figures 6B shows that among workers of prime working age, the lowest Employment-to-Population Ratios belong to Black males at 79% and Hispanic females at 70% – a 9 percentage point spread within a single age band that is otherwise the labor market's highest-participation cohort. The fact that both figures sit well below the prime-age averages for White and Asian workers of the same gender indicates that racial differences remain strongly associated with employment variation within the same age cohort. The 70% Employment-to-Population Ratio recorded for Hispanic females is the single lowest data point plotted across the prime-age figures (Table 6B). The two-proportion z-test gave z=23.71 ($p<0.001$), strongly confirming the 10.1pp gap as the sharpest prime-age disparity.

4. Occupational concentration by race produces three wholly distinct AI Exposure profiles across Hispanic, Black, and Asian workers. Figures 6C shows that the dominant occupation for each of these three groups is entirely different: in farming, fishing, and construction and extraction roles for Hispanic workers; healthcare support for the Black community; and computer and mathematics for

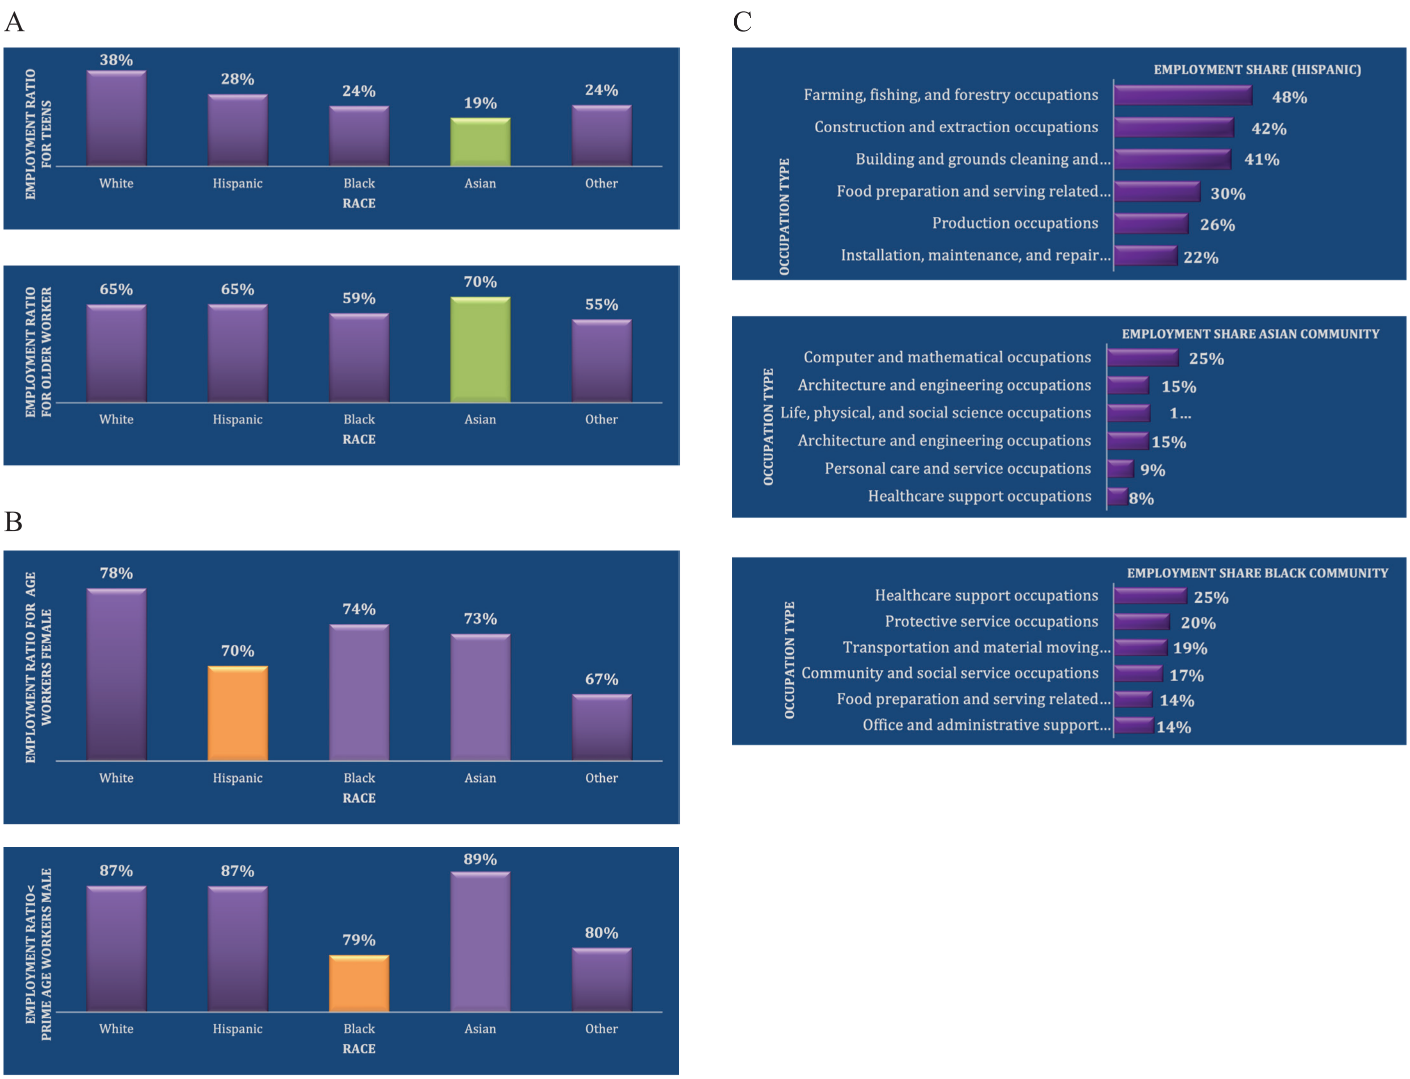

Figure 6. *EPR (%) by race/ethnicity (five categories: White, Hispanic or Latino, Black or African American, Asian, and Other [American Indian/Alaska Native, Native Hawaiian/Other Pacific Islander, and Multiracial combined]), 2022–2025: (A) youth vs. older workers; (B) primeage by sex; (C) occupational distribution. Source: IPUMS CPS, weighted.*

Asian workers. These three occupational clusters sit at opposing ends of the AI Exposure spectrum – farming and fishing carry low AI Displacement risk, healthcare support occupations register similarly low exposure, while computing and mathematics roles carry some of the highest AI Exposure scores in the dataset.

Employment Trend by Age

Understanding employment trends by age is essential because workers at different life stages face fundamentally different labor market vulnerabilities and opportunities. Age also determines occupational concentration, benefit access, and the severity of potential employment disruptions – making it a critical lens through which the potentially uneven distributional effects of market restructuring occurring during a period of AI adoption become visible.

<u>Key Findings: Employment Trends by Age</u>

1. Prime-age workers (25–54) have the highest EPR. Figures 7A and 7B show prime-age EPRs consistently above 78% throughout 2022–2025. The shift is small but statistically significant. The two-proportion z-test gave $z=-3.34$ ($p<0.001$), showing prime-age EPR rose significantly. Within this band, the 35–44 cohort carries the highest AI Exposure (rising from 0.50 to 0.63) and the largest share of high Exposure workers (39% by 2025).
2. Older worker participation (55-64) is rising and coincides with shifts in occupational mix. Figure 7 records the 65+ EPR climbing from 19.2% to 20.9%, while the 55–64 EPR rose from 60.3% to 61.2% over the same period. The two-proportion z-test gave $z=0.04$ ($p=0.97$), showing no change in 65–69 EPR and is not statistically significant. As shown in Figure 7, older workers returning to the labor market cluster in part-time service and sales/office support roles. The two-proportion z-test gave $z=-2.61$ ($p=0.009$), showing the genuine EPR rise belongs to the 55–64 band.
3. Youth labor market attachment is weakening. Figures 7A and 7B show the 16–24 EPR flat or marginally declining across the period. Workers in this cohort are disproportionately concentrated in service occupations (21.3%), including retail, food service, and hospitality. The two-proportion z-test gave $z=2.11$ ($p=0.035$), confirming the youth EPR decline as significant, supporting the claim.
4. Occupational AI Task Exposure Scores show a mixed association with age. The continuous Index (0–1) rises from 0.40 at ages 16–24 to a peak of 0.63 at 35–44, then declines – an inverted-U. A quadratic trend test on the only available field, a categorical (1–3) exposure score, did not confirm this shape ($t=-9.73$, $p<0.001$ against linearity), instead showing a monotonic rise (1.42→1.56) with no prime-age peak. The inverted-U claim should be treated as provisional pending verification with the continuous score. The 55–64 cohort, despite lower continuous-index scores, shows exposure rising to 54% by 2025 – consistent with the categorical field's monotonic pattern.

A

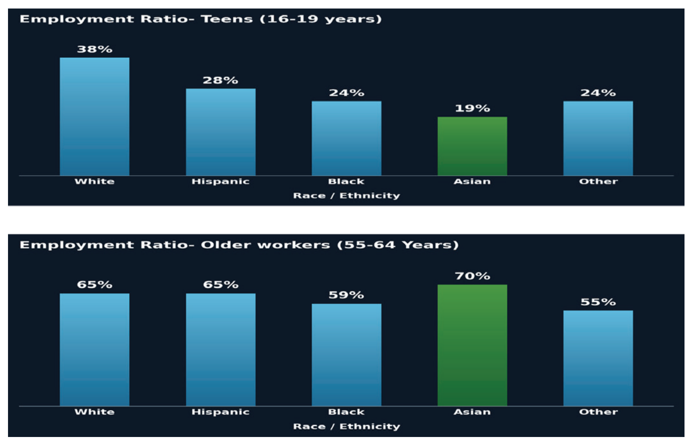

B

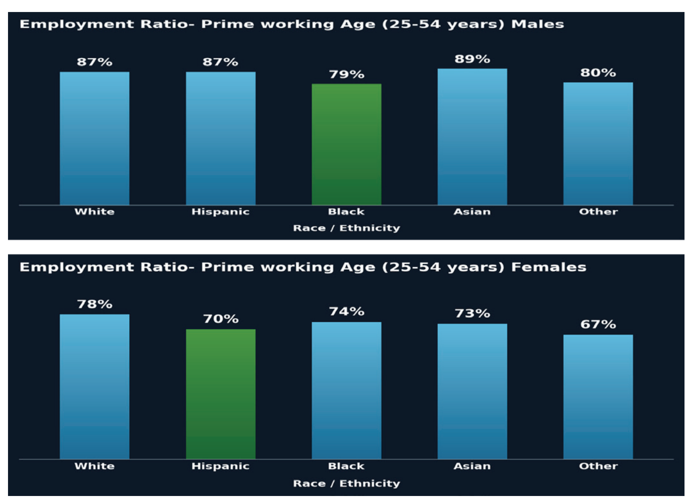

Figure 7. *EPR (%) by Prime Working Age 2022–2025: (A) by race/ethnicity (White, Hispanic or Latino, Black or African American, Asian, and Other [AIAN, NHPI, and Multiracial combined]); (B) by sex. Source: IPUMS CPS.*

3.4 BLS employment projections for the next 10 years, 2024–2034 BLS 10-year projections reflect technological, demographic, industrial, and economic factors collectively and should not be interpreted as AI-specific effect estimates. In Q3'25, BLS released its US employment projections for the next 10 years at the SOC occupation level. These employment projections estimate specific values for projected employment levels and growth rates. These projections are driven by research on factors that are expected to affect employment, which may not be reflected in historical data, such as emerging technologies and new legislation.

Adjustments based on this research are generally applied conservatively, such as when there is convincing evidence for a long-term change in employment levels. The precision in the data does not account for the inherent uncertainty of predicting long-term changes in the labor market.

Focusing on the direction and relative size of projected changes, rather than on precise point estimates, may yield similar insights into employment trends and themes across occupations and industries.

A crosswalk between the AI Exposure index at the SOC occupation level and the BLS employment projections (which includes education level required for occupations) enables insights into how AI Exposure patterns relate to projected growth and declines over the next 10 years

The BLS AI Exposure Index analysis provides a lens through which this distributional divergence is quantifiable: which occupational clusters are gaining employment share during period of AI adoption, which are contracting, and how those trajectories intersect with the credential and demographic profiles examined in preceding sections. Table 5A charts projected employment level changes from 2024 to 2034 by AI Exposure tier; Table 5B shows the same projection disaggregated by educational attainment requirement, illuminating the credential dimension of labor market restructuring occurring during a period of increasing AI adoption.

Bachelor's+ Roles: Low-to-Medium AI Exposure

Occupations requiring a bachelor's degree or higher with limited AI substitution risk are projected to add 1,558,000 jobs (+8.9%). Medical & Health Services Managers lead with +23.2% growth, consistent with rising projected demand for credentialled oversight in regulated industries (Table 6A).

AI Complementary Roles: High AI Task Exposure Occupations with Positive Projected Growth

A concentrated cluster of analytical and security-focused roles thrives despite high AI Exposure, adding 280,000 jobs (+11.5%). AI tools appear associated with task augmentation in these functions – Data Scientists (+33.5%) and Information Security Analysts (+28.5%) are the fastest-growing (Table 6A).

High AI Exposure Score: Occupations with Elevated AI Exposure Score

BLS projections (noting that multiple factors beyond AI Exposure contribute to these estimates) indicate that high AI Exposure, routine-task roles are projected employment decline of 593,000 jobs (-5.2%). Data Entry Keyers face the steepest decline (-25.9%), while Office Clerks and Customer Service Representatives account for the largest absolute losses. AI Exposure Index scores exceed 60% across all occupations in this group (Table 6B).

Education × AI Exposure Matrix

The interaction between education level and AI Exposure is a meaningful indicator of employment trajectory across all occupation segments (Figure 8A).

Table 5. *(A)BLS-projected employment change (%), 2024–2034, by AI Exposure tier (No, Low, Medium, High). (B)BLS-projected employment change (%), 2024–2034, by minimum educational-attainment requirement (six tiers).*

AI Exposure Level	Percentage Employment change
No Exposure	2.8%
Low AI Exposure	3.5%
Medium AI Exposure	3.0%
High AI Exposure	-0.6%

Education Level	Projected Employment Change Percentage
Doctoral or Professional Degree	5.9%
Master's Degree	9.2%
Bachelor's Degree	5.8%
Some College or Associate Degree	0.3%
High School Diploma or Equivalent	0.5%
Less Than High School	2.3%

Table 6. *(A)Projected employment (000s), 2024 vs. 2034, for seven bachelor's+ occupations with medium-to-low AI exposure; total +1,558,000 jobs (+8.9%); (B) Projected employment (000s), 2024 vs. 2034, for five high-AI-exposure occupations with positive growth.*

Occupation Type	2024 (000s)	2034 (000s)	Change	% Change
Healthcare Practitioners & Technical	6,248	6,739	+491	7.9%
Computer & Mathematical Occupations	3,511	3,896	+386	11.0%
General & Operations Managers	3,713	3,877	+164	4.4%
Engineers	1,790	1,913	+123	6.9%
Medical & Health Services Managers	616	759	+143	23.2%
Financial Managers	869	997	+129	14.8%
Computer & Info Systems Managers	667	769	+102	15.2%
Total	17,594	19,152	+1,558	8.9%

Occupation	2024 (000s)	2034 (000s)	Change	% Change
Information Security Analysts	183	235	+52	28.5%
Data Scientists	246	328	+82	33.5%
Operations Research Analysts	112	136	+24	21.5%
Market Research & Marketing Specialists	942	1,005	+63	6.7%
Human Resources Specialists	944	1,003	+58	6.2%
Total	2,427	2,707	+280	11.5%

Key Findings

1. The high AI Exposure tier is the only tier losing jobs, with a projected employment decline of −0.6% (Table 5), while No AI Exposure grows +2.8%, Low +3.5%, and Medium +3.0%. Projected declines in employment level are concentrated in data entry, administrative coordination, and clerical processing – roles where a substantial share of task content is measurably accessible to current AI systems. The contraction appears to be concentrated in routine-task oriented occupation types, not economy-wide.
2. Low and medium exposure occupations show comparatively stronger projected growth
   Low (+3.5%) and Medium (+3.0%) AI Exposure tiers outgrow even No Exposure (+2.8%) roles (Table 5).
3. Advanced credentials capture most projected education-linked growth
   Master's degree roles lead at +9.2%, followed by Doctoral/Professional at +5.9% and Bachelor's at +5.8% (Table 5B). The gradient then collapses sharply: Some College +0.3%, High School +0.5%.
4. The middle-credential cohort shows near zero growth
   Some College (+0.3%) and High School (+0.5%) credential tiers sit at near-zero growth – concentrated in administrative and clerical roles at the intersection of high AI Exposure and projected employment decline risk.
5. Bachelor's+ roles with low-medium AI Exposure drive 1.56M new jobs
   Seven occupational groups requiring a bachelor's degree or higher with limited substitution risk add 1,558,000 jobs (+8.9%) (Table 6A). Medical & Health Services Managers lead at +23.2%, followed by IT Managers +15.2% and Financial Managers +14.8% – all exceeding the economy-wide +2.7% average. Advanced credentials combined with low automation susceptibility are associated with comparatively stronger long-term employment projections through 2034.
6. High AI Task Exposure occupations show the fastest projected growth
   Data Scientists (+33.5%), Information Security

Analysts (+28.5%), and Operations Research Analysts (+21.5%) post the highest individual growth rates in the dataset despite high AI Exposure Index scores (Table 6B). Together these roles add 280,000 jobs by 2034.

7. Routine clerical show a projected decline 593,000 jobs
   Seven high exposure, routine-task occupations lose 593,000 jobs (−5.2%) through 2034 (Table 7). Data Entry Keyers fall −25.9%; Office Clerks lose 178,000 – the largest single absolute decline. AI Exposure Index exceeds 60% across all seven.
8. Education level is associated with variations in AI Exposure outcomes at every exposure tier
   The Education × AI Exposure matrix (Figure 8A) shows Master's-degree roles growing +11.9–12.0% even under high exposure, while Some College roles turn negative at −3.2% in the same band. Sub-high-school, high exposure roles decline −4.6%. The pattern is consistent across all tiers: higher education is associated with lower AI Task Exposure scores and stronger projected employment growth. This reflects an observed association between higher educational attainment and lower measured AI task exposure, not a demonstrated protective effect of education against automation.

The projected growth/decline figures reported in this section are BLS point estimates without published confidence intervals; comparisons between AI Exposure tiers, occupations, or education levels should therefore be read as directional rather than statistically tested.

Table 7. *Projected employment (000s), 2024 vs. 2034, for seven high-AI-exposure routine occupations in decline; total −593,000 jobs (−5.2%).*

Occupation	AI Exposure	2024 (000s)	2034 (000s)	Change	% Change
Customer Service Representatives	High	2,814	2,660	-154	-5.5%
Office Clerks, General	High	2,646	2,468	-178	-6.7%
Bookkeeping & Auditing Clerks	High	1,613	1,519	-94	-5.8%
Secretaries & Admin Assistants	High	1,944	1,913	-31	-1.6%
First-line Supervisors, Retail	High	1,433	1,360	-72	-5.0%
Computer User Support Specialists	High	730	703	-27	-3.7%
Data Entry Keyers	High	142	105	-37	-25.9%
Total		11,321	10,729	-593	-5.2%

Data Modelling: Gradient Boosted Machine (GBM) Classifier

Gradient Boosting Machine (GBM) is used to model this data. The candidate predictors: occupation description, occupation grouping, education level, age group, race/ethnicity, sex, and AI Exposure category – were chosen for their theoretical relevance to labor market outcomes and availability in the IPUMS CPS repeated cross sectional data. The Race_Grouping predictor used the same fivecategory recode reported in Figures 6–7 (White, Hispanic or Latino, Black or African American, Asian, and Other), constructed from the detailed IPUMS RACE and HISPAN codes as described in [Methods §Data Aggregation]. The outcome was a binary indicator: occupation-demographic cells were labeled "Growth" or "Decline" based on year-over-year employment change, 2022–2025. The 22,715-row cross-sectional cell dataset was split 80/20 (stratified) into training (18,172) and test (4,543) sets, preserving class balance (47.5% Growth / 52.5% Decline). Hyperparameters (depth 6, 50 leaves, learning rate 0.05) were tuned via 5-fold cross-validation with early stopping, converging at iteration 82 (Table 8A). Performance was evaluated on the held-out test set using AUC, accuracy, precision, recall, and F1-score.

<u>Rationale and Dataset</u>

Classifying occupation-demographic cells as Growth or Decline (2022–2025) requires a method tolerant of high-cardinality categoricals (500+ SOC codes), non-linear interactions, and a near-balanced binary outcome. A Light GBM GBDT classifier satisfies all three without distributional assumptions and produces calibrated class

Table 8. *(A) GBM (LightGBM) classifier setup: test AUC=0.56, accuracy=54.8%, 18,172/4,543 train/test split, class balance 47.5%/52.5%, depth 6, 50 leaves, learning rate 0.05 (B) GBM test-set precision, recall, and F1 by class (Decline, Growth) and macro average (F1=0.530); based on n=4,543 held-out cells*

Metric	Value
AUC (test)	0.5619
Accuracy	54.8%
Best iteration	82
Train / Test	18,172 / 4,543
Class balance	47.5% Growth · 52.5% Decline
Hyperparameters	Depth 6 · Leaves 50 · LR 0.05

Class	Precision	Recall	F1
Decline (0)	0.555	0.707	0.622
Growth (1)	0.534	0.371	0.438
Macro avg	0.544	0.539	0.530

probabilities. SHAP values restore interpretability post-hoc. The 22,715-row cross-sectional dataset (18,172 train / 4,543 test; 47.5% Growth / 52.5% Decline) was trained with max depth 6, 50 leaves, learning rate 0.05, converging at iteration 82. The model used Python 3.11 and LightGBM 4.x, with a fixed random seed (random_state=42) for reproducibility. Categorical predictors were passed as LightGBM native categorical features; missing values were handled via LightGBM's built-in missing-value routing. Near-zero employment change (0%) was classified as Decline. The split was performed at the cell, not occupation, level, so the same occupation may appear in both sets, risking leakage and an inflated AUC. Occupation-grouped or time-based validation would give a more conservative estimate

<u>Model Performance</u>

Higher AUC (closer to 1, e.g., 0.8–1.0): strong, robust discrimination – the model reliably ranks positives above negatives, with stable, dependable predictions. Lower AUC (closer to 0.5): weak discrimination – predictions

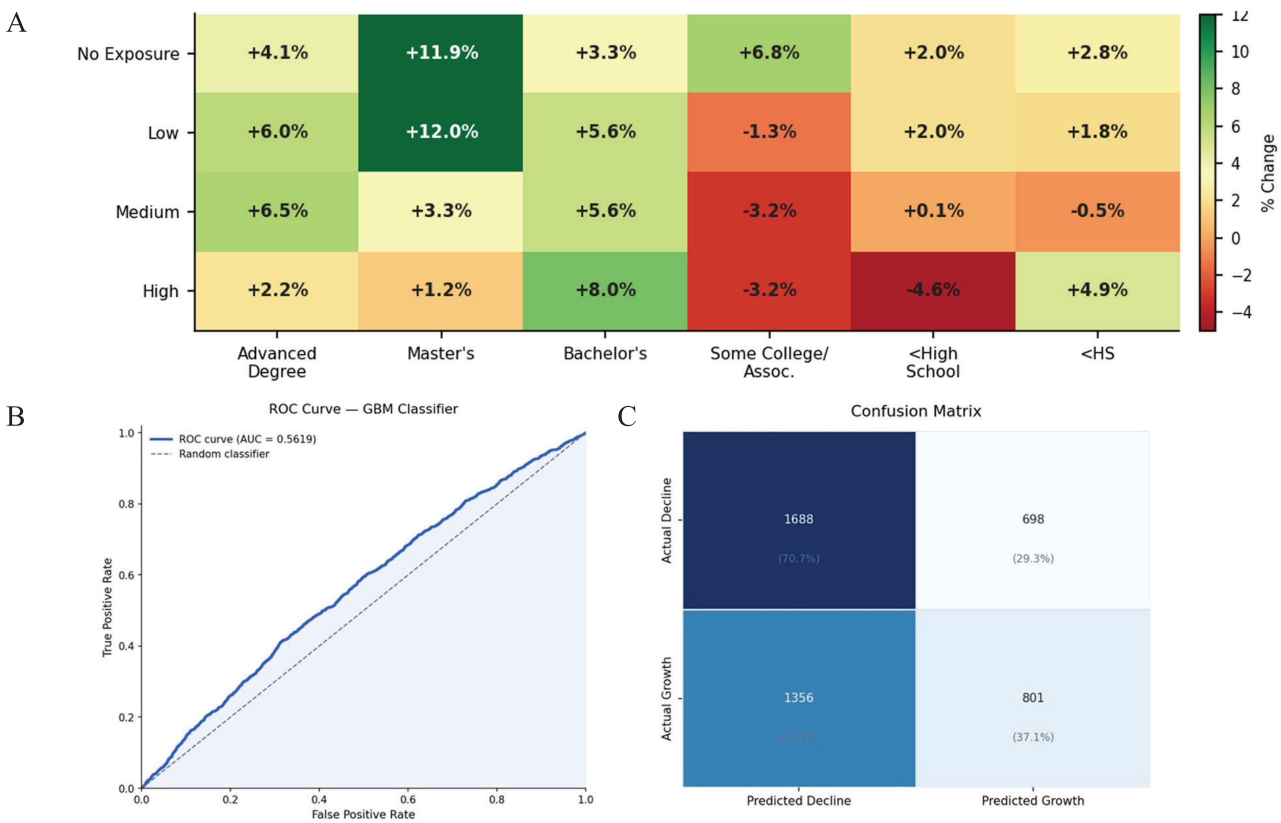

Figure 8. *(A) BLS-projected employment growth (%) by AI Exposure and education, 2024–2034; (B) GBM ROC curve (AUC=0.56); (C) confusion matrix, test set (n=4,543, Decline/Growth classes).*

are unstable and only marginally. The GBM achieved AUC 0.5619, indicating weak discrimination and only marginal improvement over chance. The model should be treated as exploratory and of limited practical predictive usefulness. The ROC curve (Figure 8B) and confusion matrix (Figure 8C) show an asymmetric error pattern: 70.7% of Decline cells identified correctly versus only 37.1% of Growth cells. Predicted probabilities cluster near 0.40–0.55 (Figure S1), reflecting the ceiling imposed by unobserved macro and firm-level factors. Table 8B reports macro F1 of 0.530.

<u>Feature Importance and SHAP Interpretability</u>

By split count (Figure S2), OCC_DESCRIPTION dominates (1,504; 44% of splits), nearly 3× Occupation_Grouping (553). By mean absolute SHAP (Figure S2), Race_Grouping leads (0.093) above OCC_DESCRIPTION (0.073), indicating race carries a higher average per-prediction impact. The beeswarm (Figure S3) confirms education's near-monotonic positive gradient toward Growth. The waterfall (Figure S4) shows the highest-confidence Growth prediction driven by OCC (+0.45), Education (+0.16), Age (+0.11), and AIExposure (+0.04). SHAP dependency plots (Figure S4) show OCC spans −1.0 to +0.4, underscoring vast within-occupation heterogeneity. As SHAP values here derive from a single model fit rather than bootstrap-refit estimates, so the ranking of Race_Grouping above OCC_DESCRIPTION should be confirmed with bootstrap-based SHAP variance before being treated as statistically robust.

<u>Feature Interactions</u>

The interaction heatmap (Figure S5) shows OCC self-interaction dominates (gain 8,746), followed by OCC × Occupation_Grouping (5,005) and Age × OCC (3,429). AIExposure × OCC (gain 1,076, rank 10) indicates AI Exposure appears to interact with occupation-level employment patterns rather than acting as a standalone predictor.

<u>Interpretation of Model results</u>

1. AUC 0.5619 – AUC of 0.56 indicates weak discrimination, only marginally above chance, signalling limited standalone predictive value for demographic, occupational, and AI Exposure variables – best treated as exploratory, not validated.
2. Confusion matrix: An asymmetric error pattern: 70.7% of Decline cells identified correctly versus only 37.1% of Growth cells. It indicates that the observable features are able to better characterize the segments that experienced highest employment decline than identifying segments that exhibited the highest job growth rate.
3. SHAP values: Race grouping has the highest SHAP value. This only indicates that race is an independent predictor not explained by other features, meaning even after controlling for other features race grouping enhances the predictability. Where 'race' refers to the fivecategory Race_Grouping variable defined in Methods (White, Hispanic or Latino, Black or African American, Asian, and Other); because AIAN, NHPI, and Multiracial respondents are pooled into 'Other,' this result cannot be attributed to any single one of those groups.
4. GBM Splits: Occupation type being the dominant feature, accounts for 44% of the splits indicating that it is the dominant predictive feature within the model.
5. AI Exposure: The low SHAP value for AI Exposure indicates that it is not an independent driver of employment change. However, its interaction with occupation indicates that AI Exposure may interact with occupational employment patterns.
6. Feature Interaction Analysis: The interactions with occupation indicate high within-occupation heterogeneity. The interactions between occupation level and age indicate that the employment trajectory within an occupation can differ by age groups.

DISCUSSION

The employment-to-population ratio held broadly stable across 2022–2025 even as average AI task-exposure scores rose economy-wide, and the classifier's predictive performance was modest, only slightly above chance. This suggests occupational and demographic structure carry some association with employment outcomes, though most variation lies outside this analysis's scope. The pattern is broadly consistent with the ILO's characterization of global employment as stable overall, and echoes evidence that aggregate indicators can mask heterogeneity during labor-market shocks (1, 2). This study extends that literature by showing heterogeneity beneath a stable aggregate ratio is patterned by AI task-exposure specifically, though the design cannot establish that exposure, rather than a correlated factor, accounts for the patterning.

Occupational composition carried the strongest association with AI exposure, with education showing a corresponding gradient; occupations concentrated in analytical, judgement-intensive work showed continued projected growth despite comparatively high exposure. This aligns with generative AI tends to augment rather than substitute for skilled task performance, and with Harvard Business School Working Knowledge's framing that AI's effect depends on whether it enhances or displaces core tasks (8, 16). One plausible reading is that higher-credential occupations involve tasks AI performs less reliably; an equally plausible alternative is that demand reflects factors unrelated to AI, such as sector growth or credential-based segmentation. The data cannot distinguish between these. The middle-credential tiers (Some College, High School) show comparatively flat growth; one plausible reading is that these roles sit between the postgraduate/analytical tier, where AI is more likely to augment judgement-intensive work, and physically-grounded occupations largely outside current AI task-exposure – though the data cannot confirm this either

AI exposure was not evenly distributed across gender and racial groups, and occupational sorting retained an association with exposure not fully explained by education or age. This is broadly consistent with Anthropic's exposure-index work, which found similar concentration nationally; this analysis adds a demographic disaggregation using CPS microdata the original index lacked (10). Whether this reflects a direct association with AI exposure or longstanding occupational segregation correlating with exposure cannot be distinguished here. One plausible explanation for the residual predictive contribution of race in the model is that Race_Grouping captures occupational clustering not fully represented by the OCC_DESCRIPTION variable; this explanation is speculative and cannot be confirmed with the present data.

The exposure-related employment differential narrowed somewhat with age, consistent with WEF and ILO observations that AI-related risk is not uniform across age groups (1, 9). Plausible, unconfirmed explanations include rising credential requirements and employer expectations for younger workers, and shorter transition windows for older workers; neither can be confirmed as the operative mechanism.

The interaction between AI exposure and occupation contributed only marginally to predictive signal, consistent with the Yale Budget Lab's assessment that detectable AI-driven disruption has so far been limited (13). Non-AI explanations – offshoring, sectoral demand shifts, macroeconomic conditions – remain equally plausible and cannot be ruled out. Relative to prior work examining exposure indices or aggregate trends alone, this study integrates occupation-level exposure, demographic microdata, and near-term projections within one framework.

These findings suggest directions that may warrant further investigation, though none has been evaluated for effectiveness here and none should be read as a demonstrated intervention: reskilling support for high-exposure, female-dominated administrative and clerical occupations; credentialing pathways for workers with some-college or high-school-only attainment; and transition support for workers aged 55–64, given the observed convergence of rising exposure and shorter transition windows. These remain tentative implications, not tested interventions.

CONCLUSION

This paper examined employment trends from 2022 to 2025 along with the BLS projections of employment levels for the next 10 years to assess the extent to which generative AI diffusion coincides with shifts in US employment patterns and which groups are disproportionately represented in occupations classified as high AI Exposure. Drawing on IPUMS CPS microdata, Anthropic's AI Exposure Index, and a GBM classifier – the analysis delivers five integrated findings.

This detailed analysis provides a framework to go beneath the headline unemployment numbers and understand, in a structured method, the associations between AI Exposure and employment patterns across demographics, skill levels, and occupation levels, spanning both augmentation and high exposure tiers. While limited in scope, such analysis is increasingly important – providing an analytical framework for new workforce entrants who will be employed over the next 2–3 decades, and may help policymakers across the economy, education, and workforce management with an assessment of economic, social, and political impacts over the next 20–30 years. Note that BLS 10-year projections should not be interpreted as AI-specific effect estimates.

LIMITATIONS AND FUTURE WORK

The impact of AI is still evolving at an accelerated pace and remains in its nascent stage; the ability to fully understand its long-term effects is therefore inherently limited. As usage and adoption of AI increases, the ability to accurately measure and assess its impact will also improve over time. The GBM model captures composition effects rather than causal displacement, and the three-year observation window may conflate AI-driven change with post-pandemic normalisation. Also, the three-year window is relatively short to fully measure impacts of potential economic cycles, geo-political disruptions and other such factors that take a longer time to manifest into employment.

Further limitations are worth noting. Firstly, the 2022–2025 window makes it difficult to separate AI-related effects from post-pandemic labor market recovery, as pandemic-era hiring unwind, remote-work normalisation, and labor-force re-entry occurred over a similar period. Secondly, the Anthropic AI Exposure Index is a proxy derived from Claude usage patterns rather than a direct measure of task-level automation and may under- or over-state exposure for occupations with low AI adoption, non-digital work, or fast-evolving task content. Thirdly, aggregating exposure scores to the occupation or demographic level can introduce aggregation bias, masking materially different exposure among workers within the same code. Fourthly, the split was performed at the cell rather than occupation level, so the same occupation may appear in both sets, risking information leakage that occupation-grouped or time-based validation would address. Finally, the analysis is observational and associational, identifying statistical relationships rather than causal effects, and cannot rule out confounding, reverse causality, or omitted variables.

Future work should incorporate panel fixed effects, the continuous AI Exposure score, multi-class classification, and occupation-grouped or time-based cross-validation to mitigate leakage from the current row-level split. To more accurately establish causality and address collinearity, more sophisticated modelling constructs should be further evaluated.

CONFLICT OF INTEREST

The authors declare that there are no conflicts of interest related to this work.

REFERENCES

1. International Labor Organization. World employment and social outlook: Trends 2025 [Internet]. Geneva: International Labour Organization; 2025. Available from: https://www.ilo.org/publications/flagship-reports/world-employment-and-social-outlook-trends-2025 (accessed on 2026-07-31).
2. Johnstone S, Saridakis G, Wilkinson A. The global financial crisis, work and employment: Ten years on. *Econ Ind Democr*. 2019; 40 (3): 455–468. https://doi.org/10.1177/0143831X19866532
3. Flood S, King M, Rodgers R, Ruggles S, Warren JR, Backman D, Breton E, Cooper G, Rivera Drew JA, Richards S, Van Riper D. Integrated Public Use Microdata Series, Current Population Survey: Version 13.0 [dataset]. Minneapolis (MN): IPUMS; 2025. https://doi.org/10.18128/D030.V13.0. Available from: https://cps.ipums.org/cps/ (accessed on 2026-07-31).
4. OpenAI. Introducing ChatGPT [Internet]. San Francisco (CA): OpenAI; 2022 Nov 30. Available from: https://openai.com/index/chatgpt/ (accessed on 2026-07-31).
5. Donovan SA, Labonte M. Labor market patterns since 2007. Washington (DC): Congressional Research Service; 2018 Oct 3. Report No.: R45330. Available from: https://www.congress.gov/crs-product/R45330 (accessed on 2026-07-31).
6. Challenger, Gray & Christmas. 2025 Year-End Challenger Report: Highest Q4 Layoffs Since 2008; Lowest YTD Hiring Since 2010 [Internet]. Chicago (IL): Challenger, Gray & Christmas; 2026 Jan 8. Available from: https://www.challengergray.com/blog/2025-year-end-challenger-report-highest-q4-layoffs-since-2008-lowest-ytd-hiring-since-2010/ (accessed on 2026-07-31).
7. CNBC. AI job cuts: Amazon, Microsoft and more cite AI for 2025 layoffs [Internet]. Englewood Cliffs (NJ): CNBC; 2025 Dec 21. Available from: https://www.cnbc.com/2025/12/21/ai-job-cuts-amazon-microsoft-and-more-cite-ai-for-2025-layoffs.html (accessed on 2026-08-02).
8. Brynjolfsson E, Li D, Raymond LR. Generative AI at work. NBER Working Paper No. 31161. 2023 Apr (revised 2024 Nov); published as Q J Econ. 2025; 140 (2): 889–944. https://doi.org/10.1093/qje/qjae044. Available from: https://www.nber.org/system/files/working_papers/w31161/w31161.pdf (accessed on 2026-07-31).
9. World Economic Forum. The jobs most likely to be lost and created because of AI [Internet]. Geneva: World Economic Forum; 2023 May. Available from: https://www.weforum.org/stories/2023/05/jobs-lost-created-ai-gpt/ (accessed on 2026-07-31).
10. Massenkoff M, McCrory P. Labor market impacts of AI: A new measure and early evidence [Internet]. San Francisco (CA): Anthropic; 2026 Mar 5. Available from: https://www.anthropic.com/research/labor-market-

impacts (accessed on 2026-07-31).
11. U.S. Bureau of Labor Statistics, U.S. Census Bureau. Monthly Current Population Survey: Public use microdata files [Internet]. Washington (DC): U.S. Census Bureau; 2024. Available from: https://www2.census.gov/programs-surveys/cps/methodology/PublicUseDocumentation_final.pdf (accessed on 2026-08-09).
12. U.S. Bureau of Labor Statistics. Table 3. Employment status of the civilian noninstitutional population by age, sex, and race, 2025 annual averages [Internet]. Washington (DC): U.S. Bureau of Labor Statistics; 2026. Available from: https://www.bls.gov/cps/cpsaat03.htm (accessed on 2026-07-31).
13. Gimbel M, Kinder M, Kendall J, Lee M. Evaluating the impact of AI on the labor market: Current state of affairs [Internet]. New Haven (CT): The Budget Lab at Yale; 2025 Oct 1. Available from: https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs (accessed on 2026-07-31).
14. U.S. Bureau of Labor Statistics, U.S. Census Bureau. Current Population Survey: Design and methodology. Technical Paper 77. 2019 Oct. Available from: https://www2.census.gov/programs-surveys/cps/methodology/CPS-Tech-Paper-77.pdf (accessed on 2026-07-31).
15. National Center for O*NET Development. O*NET 28.0 database [dataset]. Raleigh (NC): National Center for O*NET Development; 2023 Sep 5. Available from: https://www.onetcenter.org/db_releases.html (accessed on 2026-07-31).
16. Azpúrua AE. Enhance or eliminate? How AI will likely change these jobs [Internet]. Boston (MA): Harvard Business School Working Knowledge; 2026 Feb 20. Available from: https://www.library.hbs.edu/working-knowledge/enhance-or-eliminate-how-ai-will-likely-change-these-jobs (accessed on 2026-07-31).

SUPPLEMENTARY FIGURES

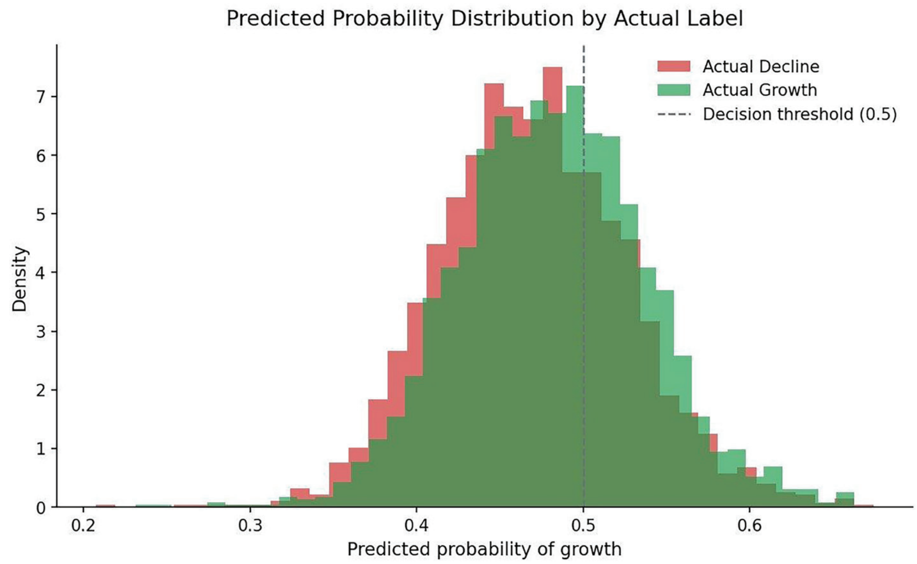

Figure S1. *Distribution of GBM-predicted "Growth" probabilities by true class, test set (n=4,543); predicted values cluster narrowly between 0.40 and 0.55.*

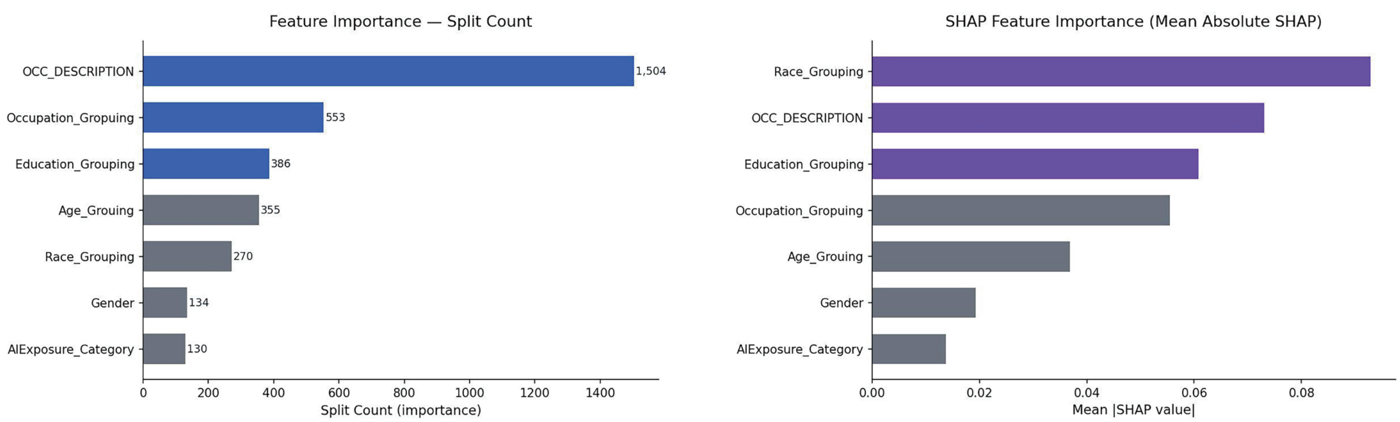

Figure S2. *(A) GBM feature importance across seven predictors, by split count and mean SHAP value; (B)occupation description dominates splits, race grouping leads mean SHAP.*

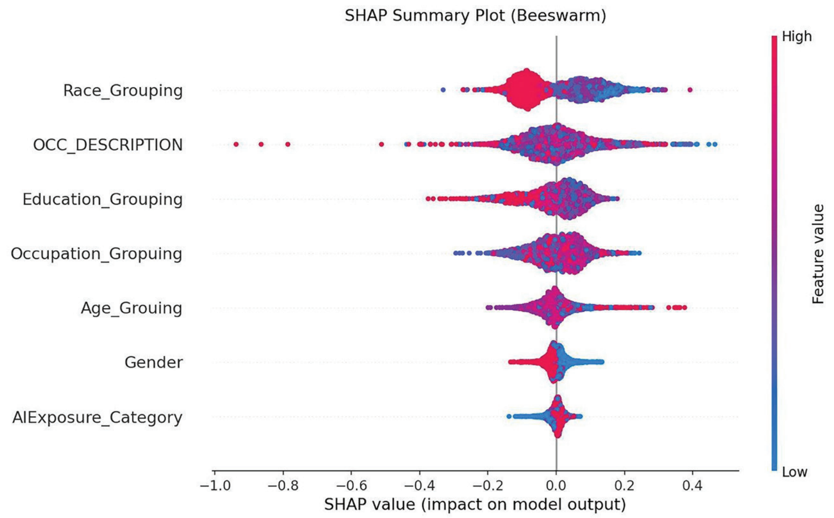

Figure S3. *SHAP beeswarm plot showing direction and magnitude of each predictor's effect on GBM Growth/Decline predictions across the test set (n=4,543).*

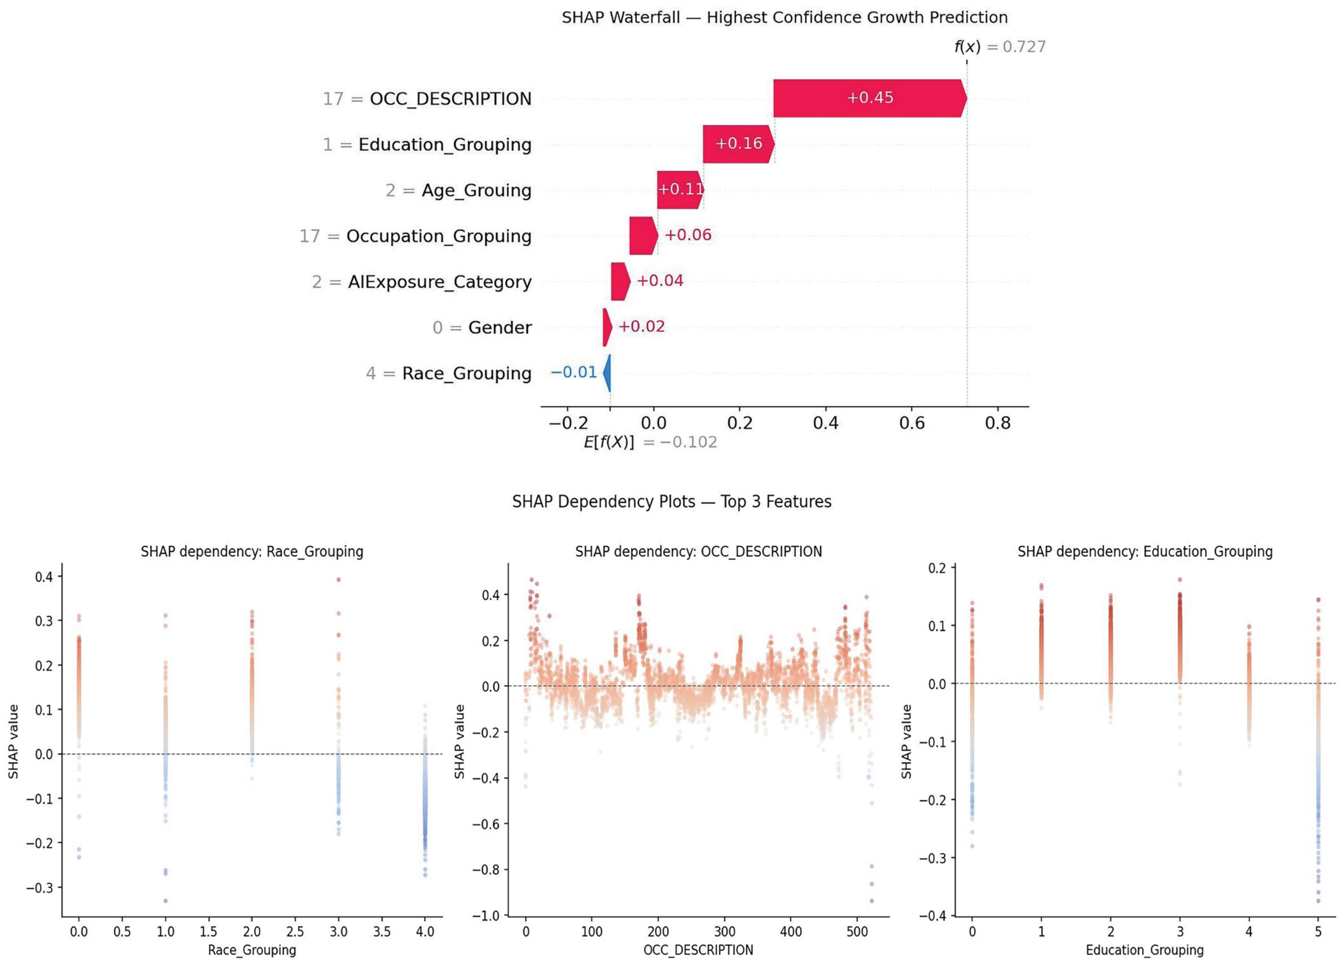

Figure S4. *SHAP waterfall for the top "Growth" prediction (occupation, education, age, AI Exposure contributions) and dependency plots for race, occupation, and education.*

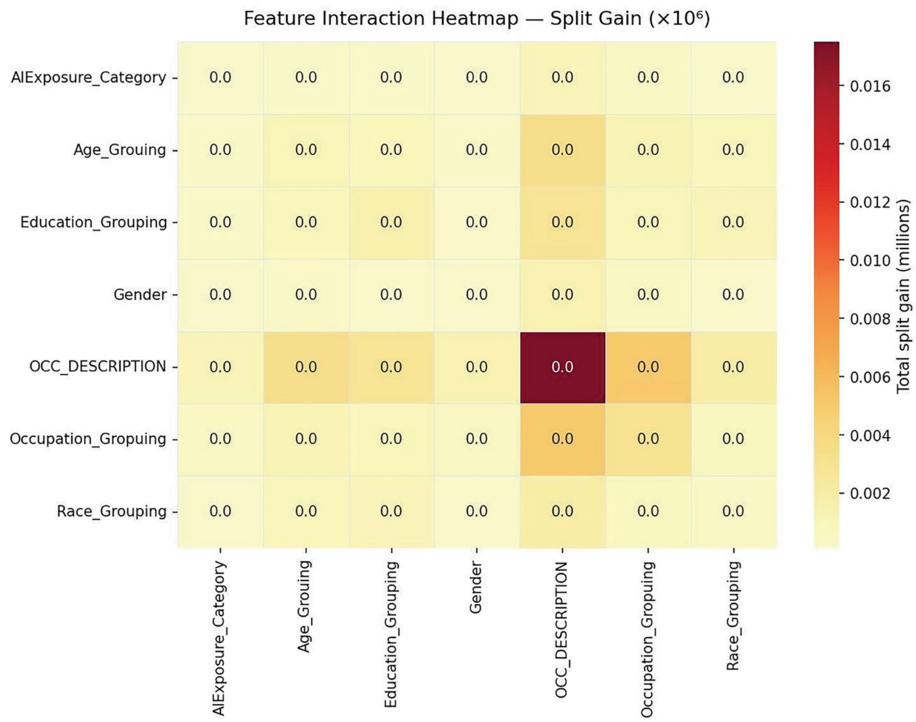

Figure S5. *Heatmap of pairwise feature-interaction strength (mean split gain, ×10⁶) among the seven GBM predictors used in the classifier.*