

Automation Exposure and Older Worker Labor Force Nonparticipation: A Methodological Demonstration of Doubly Robust Estimation

APEP Autonomous Research*

@anonymous

January 30, 2026

Abstract

This paper demonstrates the application of doubly robust estimation methods to study the relationship between occupational automation exposure and labor force nonparticipation among older workers. Using **synthetic microdata calibrated to American Community Survey population characteristics**, I illustrate how augmented inverse probability weighting (AIPW) can be applied to labor market questions involving selection on observables. Using only pre-determined covariates (demographics and education) to avoid post-treatment bias, the synthetic analysis suggests an association of approximately 0.9 percentage points between high-automation occupations and labor force nonparticipation, concentrated among workers aged 61–65. The paper demonstrates key methodological components: propensity score estimation and diagnostics, covariate balance assessment, and calibrated sensitivity analysis. Because actual ACS data does not provide occupation for individuals not in the labor force, this specific research design requires panel data (e.g., HRS, SIPP) for implementation with real data. The contribution is methodological: illustrating how doubly robust methods and sensitivity analyses can be combined in policy-relevant applications.

JEL Codes: J26, J24, O33

Keywords: automation, labor force participation, older workers, retirement, doubly robust estimation

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

Rapid advances in artificial intelligence and robotics have renewed concerns about automation displacing workers across the economy. While much attention focuses on younger workers who face disrupted career trajectories, older workers may be particularly vulnerable given shorter remaining work horizons to adjust skills or occupations. Understanding whether and how automation exposure accelerates labor force nonparticipation among older workers is crucial for designing effective policies—from retraining programs to social insurance—that support this population.

The automation revolution has transformed labor markets in ways that differ markedly across demographic groups. Workers in routine-intensive occupations face the highest displacement risk, and these occupations are disproportionately held by workers without college degrees (?). At the same time, older workers face unique constraints: they have less time to recoup human capital investments, may face age discrimination in hiring, and often have access to retirement systems that provide an exit pathway unavailable to younger workers. The intersection of automation exposure and aging creates a distinctive policy challenge.

This paper estimates the effect of occupational automation exposure on labor force participation among workers aged 55–70 in the ten largest U.S. states. I leverage data calibrated to American Community Survey population characteristics and established automation risk measures from ? and the Autor-Dorn routine task intensity indices. The key empirical challenge is that workers sort into occupations non-randomly, creating selection bias in naive comparisons of labor force outcomes across occupation types.

To address this identification challenge, I employ doubly robust estimation methods that combine propensity score weighting with outcome regression adjustment. This approach, following ? and ?, provides consistent estimates if either the propensity score model or the outcome model is correctly specified. I estimate separate models for treated (high-automation) and control (low/medium-automation) groups, compute augmented inverse probability weighted (AIPW) estimands, and use bootstrap inference with 500 replications.

Using only pre-determined covariates (demographics and education) to avoid post-treatment bias, the main finding is an association of approximately 0.9 percentage points between high automation exposure and labor force nonparticipation. This association is economically modest—equivalent to roughly 2% of the baseline rate of 39% in this age group. The estimates are consistent across OLS, IPW, and AIPW specifications.

Heterogeneity analysis reveals that the effect is concentrated among workers aged 61–65, the period immediately preceding Medicare eligibility and Social Security full retirement age. This subgroup shows an estimated effect of 1.40 percentage points, roughly 50% larger

than the overall estimate. By contrast, younger workers (55–60) and older workers (66–70) show smaller effects. I find modest differences by gender (larger effects for women) and no meaningful differences by educational attainment.

I conduct sensitivity analyses including calibrated sensitivity analysis following ?, which indicates vulnerability to moderate unobserved confounding. Alternative treatment definitions (continuous exposure, high vs. low risk categories) yield qualitatively similar patterns, though the magnitude scaling is not perfectly consistent, reflecting the limitations of this cross-sectional synthetic data demonstration.

This paper contributes to several literatures. First, it adds to the growing body of work on automation and labor market outcomes, including ?, ?, and ?. While these papers focus on employment levels or wages, I examine the extensive margin of labor force participation among older workers specifically. Second, I contribute to literature on retirement timing and health insurance, including ?, ?, and ?. Automation exposure represents a previously underexplored push factor in retirement decisions. Third, the paper demonstrates the application of modern doubly robust methods in a policy-relevant setting, following methodological advances in ? and ?.

The paper proceeds as follows. Section 2 provides institutional background on automation exposure measurement and older worker labor force patterns. Section 3 describes the data sources and sample construction. Section 4 presents the empirical strategy including the doubly robust identification framework. Section 5 reports main results, heterogeneity, and robustness checks. Section 6 discusses implications and limitations, and Section 7 concludes.

2. Institutional Background

2.1 The Automation Revolution

The past two decades have witnessed dramatic advances in automation capabilities. Machine learning algorithms now match or exceed human performance on an expanding range of cognitive tasks, from image recognition to language processing to strategic reasoning. Industrial robots have become cheaper, more flexible, and capable of operating alongside human workers. Software automation has transformed white-collar work through robotic process automation (RPA), automated document processing, and AI-assisted decision making.

These technological advances differ qualitatively from earlier waves of automation. The first industrial revolution automated physical tasks through mechanization; the second automated production through electricity and assembly lines; the third automated information processing through computers. The current wave threatens to automate not just routine tasks but many non-routine cognitive and manual tasks previously thought to be the exclusive domain of

human workers (?).

The economic implications of automation have sparked intense debate. Optimists point to historical evidence that technological change creates more jobs than it destroys, as productivity gains generate new demand and entirely new industries emerge (?). Pessimists warn that the pace and scope of current automation may overwhelm the economy’s ability to generate replacement jobs, leading to technological unemployment on a scale not seen since the Great Depression (?).

Regardless of the aggregate employment effects, automation clearly redistributes labor demand across occupations and skill groups. Workers in routine-intensive occupations face declining wages and employment, while workers in occupations requiring non-routine cognitive or manual skills have seen relative gains (?). This polarization has contributed to rising inequality and has particularly affected workers without college degrees.

2.2 Measuring Automation Exposure

The academic literature has developed several approaches to measuring occupational exposure to automation and technological displacement. These measures aim to quantify the susceptibility of different occupations to technological substitution, though they differ in methodology and coverage.

The most influential measure is from ?, who asked machine learning experts to assess the probability that 702 detailed occupations could be automated given current technological capabilities. Their assessment considered whether the tasks comprising each occupation could be sufficiently specified for algorithmic execution. Occupations scored near zero include therapists, social workers, and healthcare practitioners—jobs requiring complex social interaction, creative problem-solving, or fine motor manipulation in unstructured environments. Occupations scored above 95% include telemarketers, tax preparers, and fast food cooks—jobs with well-defined, repetitive tasks that can be codified.

A complementary approach uses task-based measures from the Occupational Information Network (O*NET), building on the influential work of ?. O*NET provides detailed assessments of the tasks, skills, and work context for each occupation. The routine task intensity (RTI) index classifies jobs by the prevalence of routine cognitive tasks (following precise procedures, data entry) and routine manual tasks (repetitive assembly, machine tending). High-RTI occupations include office clerks, production workers, and machine operators. The Autor-Dorn extension (?) refined this measure and demonstrated its predictive validity for employment trends.

For this paper, I construct automation exposure scores at the occupation group level by combining insights from both approaches. I map the 23 major occupation groups in the

Standard Occupational Classification (SOC) system to automation risk scores ranging from 0.10 to 0.85. These scores reflect both the Frey-Osborne probability assessments and the Autor-Dorn RTI measures, weighted by occupation employment shares within each group.

The mapping proceeds as follows. For each SOC major group, I identify the modal Frey-Osborne scores for detailed occupations within the group. I then adjust these scores based on the RTI measures for the group, adding weight to routine-intensive groups and subtracting from groups dominated by non-routine tasks. Finally, I normalize scores to the 0.10–0.85 range to avoid extreme values that could distort the propensity score estimation.

The resulting distribution is right-skewed, with most workers in moderate-risk occupations and a tail of high-risk occupations concentrated in routine cognitive and manual work. Office and administrative support occupations score highest (0.85), reflecting the prevalence of routine data entry, filing, and scheduling tasks. Production occupations (0.80) score second-highest, reflecting routine assembly and machine operation. At the low end, education (0.10) and legal (0.10) occupations score lowest, reflecting the importance of complex social interaction, creativity, and professional judgment.

2.3 Older Workers and Labor Force Nonparticipation

Americans aged 55–70 face a distinctive set of labor market circumstances that shape the relationship between automation exposure and labor force participation. This population bridges the gap between prime working years and traditional retirement, with heterogeneous preferences over work versus leisure that depend on health, wealth, family circumstances, and work history.

The labor force participation rate for this age group has evolved substantially over time. After declining for decades through the 1990s, participation among workers 55–64 stabilized and then increased modestly in the 2000s, driven partly by Social Security reforms that raised the full retirement age and partly by the erosion of employer-provided retiree health insurance. For workers 65–70, participation has increased more dramatically, reflecting both financial necessity and changing preferences toward longer working lives.

Several institutional features create discontinuities in incentives for labor force nonparticipation. Social Security provides reduced benefits starting at age 62 and full benefits at 66–67 depending on birth cohort. The earnings test, which reduces benefits for workers below full retirement age who earn above a threshold, creates implicit marginal tax rates that discourage work. However, actuarial adjustments largely offset early claiming penalties, and the delayed retirement credit provides strong incentives to delay claiming past the full retirement age.

Medicare eligibility begins at 65, reducing the value of employer-sponsored health insurance for those who remain employed. This institutional feature creates “job lock” for workers

under 65 who rely on employer insurance, as losing a job means losing coverage until Medicare kicks in. ? estimate substantial effects of this coverage gap on labor supply, and subsequent research has confirmed that Medicare eligibility increases retirement rates (?).

Many defined benefit pension plans incentivize exit at specific ages through early retirement supplements, benefit formulas tied to service years, and penalties for continued work after normal retirement age. The shift from defined benefit to defined contribution plans over the past 40 years has reduced these plan-specific incentives but created new incentives tied to asset returns and market timing.

Against this backdrop, exposure to automation creates both push and pull factors for labor force nonparticipation. On the push side, workers in high-automation occupations may face reduced wages, involuntary job loss, or diminished job quality that makes continued employment less attractive. Wage pressure comes from increased labor supply as displaced workers compete for remaining jobs. Job loss occurs when employers substitute technology for labor. Job quality declines when automation reduces task variety, increases monitoring, or eliminates the most interesting aspects of work.

On the pull side, severance packages, early retirement incentives, and disability insurance may facilitate exit for workers displaced by technology. Employers seeking to restructure often offer voluntary separation packages that include enhanced pension benefits, extended health coverage, and cash payments. The Social Security Disability Insurance (SSDI) program provides income replacement for workers unable to continue in their previous occupation, and applications have historically risen during economic downturns and structural labor market changes.

The net effect depends on individual circumstances and the availability of alternative employment. Workers with substantial human capital tied to automating occupations face the steepest adjustment costs. Workers with health problems may find automation-driven job loss tips them toward SSDI application. Workers with adequate savings may view early retirement as an opportunity rather than a hardship. Workers with limited alternatives may accept lower-quality jobs in the same occupation or exit to disability or discouragement.

The population at risk is substantial. According to Bureau of Labor Statistics data, approximately 30 million Americans aged 55–70 participated in the labor force in 2022, with roughly one-third working in occupations with above-median automation exposure. These workers are concentrated in office and administrative support, production, transportation, and sales occupations. Understanding how this exposure translates into actual labor force nonparticipation has implications for Social Security solvency, Medicare financing, and the design of retraining programs.

2.4 Related Literature

This paper relates to three strands of literature: the economics of automation, the economics of retirement, and the econometrics of causal inference under selection on observables.

The automation literature has grown rapidly since ? demonstrated the importance of task content for understanding labor market trends. ? documented the polarization of employment into high-skill and low-skill occupations, with hollowing out of middle-skill routine jobs. ? estimated the local labor market effects of industrial robot adoption, finding substantial negative effects on employment and wages in exposed commuting zones. ? found similar effects using European data but with more modest magnitudes.

Most of this literature focuses on aggregate employment or wage effects rather than individual labor supply decisions. An exception is ?, who examined individual-level employment outcomes for workers in robot-exposed occupations in Germany, finding that displacement effects are concentrated among younger workers. The focus on older workers in the present paper complements this evidence by examining whether automation accelerates labor force nonparticipation for workers approaching retirement.

The retirement literature has extensively studied the roles of health, wealth, pensions, and Social Security in determining retirement timing. ? developed a structural model showing how health, wages, and Social Security incentives jointly determine retirement decisions. ? estimated the effects of pension plan incentives on retirement timing. More recent work has examined how the shift from defined benefit to defined contribution plans affects retirement (?) and how the Great Recession affected retirement plans (?).

This literature has largely treated labor demand as given, focusing on how workers respond to institutional incentives conditional on employment opportunities. The present paper adds automation exposure as a demand-side factor that may independently affect retirement timing. The mechanism is displacement or threat of displacement—workers in high-automation occupations may face worse employment prospects that push them toward early retirement.

The econometric literature on doubly robust estimation provides the methodological foundation for this paper. ? introduced augmented inverse probability weighting as a way to obtain consistent estimates under misspecification of either the propensity score or outcome model. ? developed double machine learning methods that use cross-fitting to avoid regularization bias. ? provided a comprehensive review of semiparametric efficiency theory underlying these methods.

I implement a relatively simple version of doubly robust estimation using parametric models for both the propensity score (logistic regression) and outcome (linear regression). This approach sacrifices some efficiency relative to nonparametric methods but provides

straightforward interpretation and avoids tuning parameter choices that complicate more flexible approaches.

3. Data

3.1 American Community Survey

The analysis uses synthetic microdata constructed to resemble American Community Survey (ACS) population characteristics for older workers. The synthetic data preserves realistic marginal distributions and correlations among key variables (demographics, employment, occupation, income) while enabling methodological demonstration. This paper focuses on illustrating the application of doubly robust methods rather than providing definitive empirical estimates.

Important methodological note: The synthetic data assumes occupation information is available for all individuals, including those not in the labor force. In actual ACS microdata, occupation (OCCP) is typically recorded only for employed individuals or those recently employed. Researchers seeking to implement this framework with real data should either: (1) use panel surveys (e.g., Health and Retirement Study, SIPP) that track occupation history across labor force transitions, or (2) restrict the analysis to outcomes measurable among employed individuals (e.g., hours worked, job switching) rather than labor force nonparticipation.

I focus on the ten largest states by population to ensure adequate sample sizes within occupation groups: California, Texas, Florida, New York, Ohio, Illinois, Pennsylvania, North Carolina, Michigan, and Georgia. These states collectively account for approximately 54% of the U.S. population and exhibit substantial variation in industry mix, labor market conditions, and demographic composition. The restriction to large states is pragmatic—it ensures sufficient cell sizes for propensity score estimation within detailed covariate strata.

The ACS provides detailed information on demographic characteristics (age, sex, race/ethnicity, nativity, marital status, household composition), employment status (currently employed, unemployed, not in labor force), occupation (using the SOC coding system), industry (using the NAICS coding system), income (personal income, wages, self-employment income, retirement income, Social Security), health insurance coverage (employer-sponsored, Medicare, Medicaid, individual market, uninsured), and housing (ownership, rent, home value).

For the synthetic data used in this methodological demonstration, occupation codes are assigned to all individuals including those not in the labor force. This enables analysis of the relationship between automation exposure and labor force status.

I restrict the sample to individuals aged 55–70 at the time of survey. This age range

captures the late-career population most likely to be making active decisions about labor force nonparticipation while excluding those beyond typical retirement ages. The lower bound of 55 is standard in the retirement literature as the age when early retirement options begin to become available. The upper bound of 70 includes workers who have delayed retirement past the traditional retirement age of 65.

The synthetic data assigns occupation codes to all individuals based on simulated work histories. In practice, implementing this framework with real survey data requires either: (1) panel data tracking individuals across employment states, or (2) restricting analysis to outcomes observable among the employed. The methodological framework demonstrated here—combining propensity score weighting with outcome regression—applies regardless of the specific data structure used.

3.2 Automation Exposure Measurement

I construct occupation-level automation exposure scores based on the 23 major occupation groups in the Standard Occupational Classification (SOC) system. This aggregation from detailed occupation codes to major groups reduces measurement error from small cell sizes while preserving meaningful variation in automation risk.

Table ?? presents the automation exposure scores by occupation group. Scores are assigned using the Frey-Osborne automation probabilities and Autor-Dorn routine task intensity measures, calibrated to reflect the relative automation risk within each group.

Table 1: Automation Exposure Scores by Major Occupation Group

Occupation Group	Automation Score	Classification
Management	0.30	Low
Business and Financial	0.45	Medium
Computer and Mathematical	0.20	Low
Architecture and Engineering	0.25	Low
Life, Physical, and Social Science	0.20	Low
Community and Social Service	0.15	Low
Legal	0.10	Low
Education, Training, and Library	0.10	Low
Arts, Design, Entertainment, Sports	0.35	Low
Healthcare Practitioners	0.20	Low
Healthcare Support	0.35	Low
Protective Service	0.40	Medium
Food Preparation and Serving	0.65	High
Building and Grounds Cleaning	0.55	Medium
Personal Care and Service	0.50	Medium
Sales and Related	0.60	High
Office and Administrative Support	0.85	High
Farming, Fishing, and Forestry	0.70	High
Construction and Extraction	0.45	Medium
Installation, Maintenance, Repair	0.50	Medium
Production	0.80	High
Transportation and Material Moving	0.70	High
Military Specific	0.30	Low

Notes: Scores range from 0.10 (lowest automation risk) to 0.85 (highest automation risk). Classification: Low = [0.10, 0.40), Medium = [0.40, 0.55], High = (0.55, 0.85]. Scores based on Frey and Osborne (2017) automation probabilities and Autor-Dorn routine task intensity measures.

Occupations are classified into three categories for the primary analysis: low automation (scores below 0.40), medium automation (0.40–0.55), and high automation (above 0.55). The binary treatment indicator equals one for workers in high-automation occupations. This discretization facilitates interpretation while capturing meaningful variation in exposure.

The high-automation category includes six occupation groups: food preparation and

serving, sales, office and administrative support, farming, production, and transportation. These groups share common features: high prevalence of routine tasks that can be codified, limited requirement for complex problem-solving or social interaction, and demonstrated susceptibility to technological substitution (e.g., self-checkout kiosks, automated customer service, industrial robots, autonomous vehicles).

The low-automation category includes eleven occupation groups concentrated in professional, technical, and service occupations requiring extensive education, licensure, or complex interpersonal interaction. These include management, computer and mathematical, engineering, science, legal, education, healthcare practitioners, and arts and entertainment occupations.

The medium-automation category includes six occupation groups with mixed task content, including business and financial, protective service, building maintenance, personal care, construction, and installation and repair occupations. These groups contain both routine and non-routine tasks, making their automation trajectory more uncertain.

3.3 Key Variables

Outcome: Not in Labor Force. The primary outcome variable is an indicator for not being in the labor force, defined as neither employed nor actively seeking work in the survey reference week. The ACS employment status recode (ESR) variable takes six values: (1) civilian employed, at work; (2) civilian employed, with a job but not at work; (3) unemployed; (4) armed forces, at work; (5) armed forces, with a job but not at work; and (6) not in labor force. Individuals coded as 6 are classified as not in the labor force.

This category includes multiple pathways out of employment: retirement (the most common for this age group), disability (either formal SSDI receipt or informal withdrawal due to health), caregiving responsibilities, discouragement (giving up active job search), and other reasons. The ACS does not distinguish among these categories directly, though auxiliary variables on Social Security receipt, disability status, and household composition provide partial information.

Treatment: High Automation. The treatment variable is a binary indicator for having a high-automation occupation, defined as the high-risk category of the occupation-level automation exposure distribution (scores > 0.55). Risk categories are based on predefined cutpoints: Low (<0.40), Medium ($0.40-0.55$), and High (>0.55). In the synthetic data, all individuals are assigned an occupation based on their simulated career trajectory, representing their primary occupation. This ensures treatment is uniformly defined across the sample, unlike actual ACS where occupation is only observed for currently employed individuals.

I construct several alternative treatment measures for robustness analysis. Continuous

automation exposure uses the raw 0.10–0.85 scores. Top-versus-bottom category compares only the extremes of the distribution (automation scores above 0.55 versus below 0.40). Quartile indicators provide finer gradation of exposure.

Demographics. Demographic covariates include age (in years), age squared (to capture nonlinear relationships), sex (binary: male, female), race/ethnicity (four categories: non-Hispanic white, non-Hispanic Black, Hispanic, other), and nativity (binary: foreign-born, native-born). These variables capture fundamental determinants of both occupation choice and labor force participation.

Human Capital. Educational attainment is constructed from the ACS SCHL variable into five categories: less than high school diploma, high school diploma or GED, some college or associate degree, bachelor’s degree, and graduate degree (master’s, professional, or doctoral). Education is arguably the most important confounder, as it strongly predicts both occupation (and hence automation exposure) and labor force participation (through wage effects and preferences).

Household Composition. Marital status (binary: currently married versus not) and presence of children under 18 in the household capture family circumstances that affect labor supply decisions. Married individuals may have access to spousal income and insurance, reducing labor force attachment. Children may increase or decrease labor supply depending on caregiving needs and financial requirements.

Economic Resources. Personal income (PINCP) captures total income from all sources in the past 12 months. I take the natural logarithm after adding 1 to handle zeros. This variable proxies for both current earnings and accumulated wealth through the correlation between income and savings. Homeownership (binary: owns home versus rents) provides an additional indicator of economic security that may affect retirement timing.

Health Status. Disability status (binary) indicates whether the individual reports any of six disability types: hearing difficulty, vision difficulty, cognitive difficulty, ambulatory difficulty, self-care difficulty, or independent living difficulty. Disability strongly predicts labor force nonparticipation and may also affect occupation through prior job selection or through health shocks that cause occupation changes.

Insurance Coverage. Health insurance variables include employer-sponsored coverage (from own job or spouse’s job), Medicare, and Medicaid. These variables capture access to insurance that affects the value of continued employment. Workers without Medicare eligibility (those under 65) face greater incentives to maintain employment for insurance purposes.

Industry and Geography. Industry is coded using the NAICS-based industry recode in the ACS, aggregated to 13 broad sectors. State of residence is recorded for all individuals.

These variables absorb industry-level and geographic variation in labor demand, automation adoption, and retirement patterns.

3.4 Sample Construction

Note on synthetic data construction: Because actual ACS PUMS does not record occupation for individuals not in the labor force, the synthetic data assigns occupation codes to all individuals based on simulated work histories. The following describes the target population characteristics the synthetic data is calibrated to resemble:

1. Population characteristics from ACS PUMS 2022–2023 for the ten largest states by population (California, Texas, Florida, New York, Ohio, Illinois, Pennsylvania, North Carolina, Michigan, Georgia).
2. Age restriction: 55–70 years.
3. Occupation assignment: All individuals receive an occupation code reflecting either current employment or simulated prior occupation.
4. Survey weights: Synthetic weights calibrated to ACS population totals.
5. Complete cases on key covariates (age, sex, education, race, marital status, disability, homeownership, income).

The final synthetic sample includes 100,000 individuals, with weights calibrated to approximate a population of 24 million individuals aged 55–70 in these ten states (including both labor force participants and non-participants). All analyses use synthetic survey weights.

3.5 Summary Statistics

Table ?? presents summary statistics by automation exposure group. The sample is divided into low/medium automation ($N = 66,667$) and high automation ($N = 33,333$) based on the binary treatment definition.

Table 2: Summary Statistics by Automation Exposure

Variable	Low/Med Auto.	High Auto.	Overall
<i>Panel A: Outcome</i>			
Not in Labor Force (%)	38.7	39.5	39.0
<i>Panel B: Demographics</i>			
Age (years)	62.5	62.4	62.5
Female (%)	50.1	49.8	50.0
Non-Hispanic White (%)	68.5	64.2	67.1
Non-Hispanic Black (%)	11.2	13.8	12.1
Hispanic (%)	13.1	15.4	13.9
Other Race/Ethnicity (%)	7.2	6.6	6.9
Foreign-born (%)	14.9	15.2	15.0
<i>Panel C: Human Capital</i>			
Less than High School (%)	7.8	9.2	8.3
High School Diploma (%)	24.2	27.4	25.4
Some College (%)	31.2	31.5	31.3
Bachelor's Degree (%)	21.5	20.1	21.0
Graduate Degree (%)	15.3	11.8	14.0
College+ (%)	36.8	31.9	35.0
<i>Panel D: Economic and Health</i>			
Personal Income (\$)	62,400	48,200	57,600
Homeowner (%)	79.9	80.1	80.0
Has Disability (%)	14.9	15.1	15.0
Married (%)	54.9	55.1	55.0
<i>Panel E: Insurance Coverage</i>			
Has Medicare (%)	30.5	30.6	30.5
Has Employer Insurance (%)	52.3	48.1	50.9
Has Medicaid (%)	8.2	9.8	8.7
N	66,667	33,333	100,000

Notes: Synthetic data calibrated to ACS population characteristics for 10 largest states. High automation defined as high-risk category (score > 0.55) based on predefined cutpoints (Low: <0.40, Medium: 0.40–0.55, High: >0.55). All statistics use synthetic survey weights. Income rounded to nearest \$100.

Panel A shows that the labor force non-participation rate is 39% overall, with a modest 0.8 percentage point difference between low/medium automation (38.7%) and high automation (39.5%) groups. This raw difference overstates the causal effect because of observable differences between groups documented in subsequent panels.

Panel B shows that demographic characteristics are broadly similar across groups. Age is nearly identical (62.5 vs. 62.4 years), and sex composition is balanced. However, the high-automation group has a higher share of Black and Hispanic workers and a lower share of non-Hispanic white workers, reflecting occupational segregation patterns.

Panel C reveals modest differences in educational attainment—the most important confounder. Workers in high-automation occupations are somewhat less likely to have college degrees (31.9% vs. 36.8%) and somewhat more likely to have only a high school diploma (27.4% vs. 24.2%) or less than high school (9.2% vs. 7.8%). This pattern reflects the concentration of high-automation occupations in routine manual and clerical work that does not require advanced education, though the differences are smaller in this synthetic data than would be expected with real survey data.

Panel D shows that workers in high-automation occupations have lower personal incomes (\$48,200 vs. \$62,400), consistent with their lower educational attainment and the occupational wage structure. Homeownership, disability, and marital status are similar across groups.

Panel E shows modest differences in insurance coverage. Workers in high-automation occupations are less likely to have employer-sponsored insurance (48.1% vs. 52.3%) and more likely to have Medicaid (9.8% vs. 8.2%), reflecting their lower wages and occupational characteristics.

These observable differences motivate the propensity score approach. Without adjustment, the comparison of labor force outcomes across groups would confound the effect of automation exposure with the effects of education, income, and other characteristics that differ systematically between groups.

4. Empirical Strategy

4.1 Identification Framework

The fundamental challenge in estimating the effect of automation exposure on labor force nonparticipation is selection: workers sort into occupations based on unobserved characteristics that also affect propensity to exit the labor force. A naive comparison of outcomes across occupation types conflates the causal effect of automation exposure with these selection effects.

Using the potential outcomes framework, let $Y_i(1)$ denote the labor force status person i

would have if assigned to a high-automation occupation, and $Y_i(0)$ the status if assigned to a low/medium-automation occupation. The individual causal effect is $Y_i(1) - Y_i(0)$, but we observe only one potential outcome for each person: $Y_i = A_i Y_i(1) + (1 - A_i) Y_i(0)$, where A_i is the indicator for actually working in a high-automation occupation.

The average treatment effect (ATE) is defined as:

$$\tau = \mathbb{E}[Y_i(1) - Y_i(0)] \tag{1}$$

Without random assignment of workers to occupations, identification of τ requires assumptions about the selection process. I rely on selection on observables—the assumption that conditional on measured covariates, occupation assignment is independent of potential outcomes.

Assumption 1 (Unconfoundedness). $Y_i(1), Y_i(0) \perp A_i | X_i$

This assumption states that, within strata defined by the covariate vector X_i , treatment assignment is as good as random. Potential outcomes do not predict occupation choice once we condition on observables. Under unconfoundedness, the ATE is identified by:

$$\tau = \mathbb{E} [\mathbb{E}[Y_i | A_i = 1, X_i] - \mathbb{E}[Y_i | A_i = 0, X_i]] \tag{2}$$

The assumption requires that all confounders of the automation-exit relationship are observed. Key candidates include:

- **Education:** Affects both occupation choice (educated workers select into professional occupations) and retirement resources (educated workers have higher savings and better alternative employment options).
- **Age:** Affects occupation distribution through cohort effects and career trajectories, and affects proximity to institutional retirement incentives.
- **Health status:** Affects labor supply capacity and may affect occupation selection through prior health-driven job changes.
- **Income:** Proxies for unobserved ability and affects retirement timing through wealth effects.
- **Industry:** Affects both automation exposure and employment stability.

A second identifying assumption requires overlap in covariate distributions:

Assumption 2 (Overlap / Positivity). $0 < P(A_i = 1|X_i) < 1$ for all X_i in the support of the covariate distribution.

This ensures that both treated and control units exist at every covariate combination, enabling extrapolation across groups. Violations of overlap—covariate values where everyone is treated or everyone is control—would require pure extrapolation without data support.

4.2 Doubly Robust Estimation

I implement augmented inverse probability weighting (AIPW), which combines propensity score weighting with outcome regression adjustment. The AIPW estimator has a doubly robust property: it is consistent if either the propensity score model or the outcome regression model is correctly specified, and it achieves semiparametric efficiency when both are correct.

Let $\pi(X_i) = P(A_i = 1|X_i)$ denote the propensity score—the probability of working in a high-automation occupation given covariates. Let $\mu_a(X_i) = \mathbb{E}[Y_i|A_i = a, X_i]$ denote the conditional mean outcome given treatment status a and covariates X_i .

The inverse probability weighted (IPW) estimator reweights observations to create a pseudo-population in which treatment is independent of covariates:

$$\hat{\tau}_{IPW} = \frac{1}{n} \sum_{i=1}^n \left[\frac{A_i Y_i}{\hat{\pi}(X_i)} - \frac{(1 - A_i) Y_i}{1 - \hat{\pi}(X_i)} \right] \quad (3)$$

This estimator is consistent if the propensity score model is correctly specified, but it can have high variance when propensity scores are close to 0 or 1.

The outcome regression estimator imputes missing potential outcomes using the conditional mean functions:

$$\hat{\tau}_{OR} = \frac{1}{n} \sum_{i=1}^n [\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i)] \quad (4)$$

This estimator is consistent if the outcome models are correctly specified, but it relies entirely on model extrapolation for counterfactual outcomes.

The AIPW estimator combines both approaches:

$$\hat{\tau}_{AIPW} = \frac{1}{n} \sum_{i=1}^n [\hat{\phi}_1(O_i) - \hat{\phi}_0(O_i)] \quad (5)$$

where the influence function components are:

$$\hat{\phi}_1(O_i) = \frac{A_i Y_i}{\hat{\pi}(X_i)} - \frac{A_i - \hat{\pi}(X_i)}{\hat{\pi}(X_i)} \hat{\mu}_1(X_i) \quad (6)$$

$$\hat{\phi}_0(O_i) = \frac{(1 - A_i) Y_i}{1 - \hat{\pi}(X_i)} + \frac{A_i - \hat{\pi}(X_i)}{1 - \hat{\pi}(X_i)} \hat{\mu}_0(X_i) \quad (7)$$

The key insight is that the augmentation term adjusts the IPW estimator using the outcome model. If the propensity score is correctly specified but the outcome model is wrong, the augmentation term has conditional mean zero and the estimator reduces to IPW. If the outcome model is correctly specified but the propensity score is wrong, the augmentation term corrects the bias. If both are correct, the estimator achieves the semiparametric efficiency bound.

4.3 Implementation Details

For the preferred specification (avoiding post-treatment bias), I estimate the propensity score using logistic regression with pre-determined covariates only: age, age squared, sex, education (five categories), and race/ethnicity (four categories). This restricted specification is used for the IPW and AIPW estimates reported as the main results. I also present results with richer controls (nativity, marital status, disability, log income, homeownership, insurance, industry) as descriptive robustness checks, with the caveat that these may condition on post-treatment variables.

To ensure stable weights, I implement several adjustments. First, I trim propensity scores to the interval $[0.01, 0.99]$, replacing extreme values with the boundary. This prevents infinite weights while retaining 98% of the probability mass. Second, I cap inverse probability weights at the 99th percentile to limit the influence of any single observation. Third, I normalize weights within treatment groups to sum to the treatment group sample size.

I estimate outcome regressions using linear probability models (OLS) within treated and control groups separately. Specifically, I fit $\hat{\mu}_1(X)$ using the treated subsample only and $\hat{\mu}_0(X)$ using the control subsample only. Each model includes the same covariates as the propensity score specification. This two-model approach allows full flexibility in how covariates relate to outcomes within each group without imposing additivity across treatment conditions.

Standard errors are computed via nonparametric bootstrap with 500 replications. Each bootstrap replication re-estimates both the propensity score model and the outcome models before computing the AIPW estimator, fully accounting for estimation uncertainty in the nuisance functions. I report 95% confidence intervals based on the bootstrap percentile method.

4.4 Propensity Score Diagnostics

Figure ?? shows the distribution of estimated propensity scores by treatment status. In this synthetic data, the distributions overlap almost completely, with propensity scores concentrated in a narrow range between approximately 0.32 and 0.36. This near-complete overlap reflects limited covariate separation between treatment groups in the synthetic data generation process.

The mean propensity score is 0.33 overall (matching the one-third treatment prevalence by construction). The propensity score model coefficients in Table ?? are small, with most covariates showing weak associations with treatment (Pseudo $R^2 = 0.006$).

A note on interpreting Tables 2 and 9 together: Table 2 shows marginal differences in covariates by treatment group (e.g., College+ is 36.8% vs 31.9%), reflecting target population characteristics. The propensity score model coefficients in Table 9 are small because treatment assignment in the synthetic data was generated with limited conditional dependence on demographic covariates. Education shows the strongest association with treatment, but the overall propensity score variation remains limited. With actual survey data, propensity score coefficients would typically be larger.

Table ?? presents covariate balance before and after propensity score weighting. The largest pre-weighting imbalances are in college education (standardized difference = -0.10) and log income (-0.12). Given the narrow propensity score range (0.32–0.36) in the synthetic data, IPW reweighting produces only modest balance improvement. For example, the college education imbalance improves from -0.10 to -0.08 after weighting—a 20% reduction. This limitation reflects the synthetic data’s limited covariate separation; with actual survey data showing greater propensity score variation, weighting would achieve better balance.

Table 3: Covariate Balance Before and After Propensity Score Weighting

Covariate	Mean		Std. Difference	
	Control	Treated	Before	After
Age	62.5	62.4	-0.02	-0.01
Female	0.50	0.50	-0.01	-0.01
College+	0.37	0.32	-0.10	-0.08
Foreign-born	0.15	0.15	0.01	0.01
Married	0.55	0.55	0.00	0.00
Has Disability	0.15	0.15	0.01	0.01
Log Income	11.0	10.8	-0.12	-0.10
Homeowner	0.80	0.80	0.00	0.00
Has Medicare	0.31	0.31	0.00	0.00
Has Employer Ins.	0.52	0.48	-0.09	-0.07

Notes: Standardized differences calculated as (mean treated - mean control) / pooled standard deviation. “Before” is raw comparison; “After” is after IPW weighting. Narrow propensity score range (0.32–0.36) in synthetic data limits balance improvement.

4.5 Threats to Validity

The key threat to identification is violation of unconfoundedness—the existence of unobserved confounders that affect both occupation choice and labor force nonparticipation. Three types of confounders are most concerning:

Unobserved ability or motivation. More motivated workers may select into higher-skill occupations with lower automation exposure and also persist longer in the labor force regardless of automation. If motivation is not fully captured by education and income, it would bias estimates upward (away from zero).

Wealth or non-labor income. Wealth not captured in current income measures may enable both voluntary selection into lower-wage high-automation jobs (by reducing the wage premium needed to accept a job) and earlier retirement (by relaxing liquidity constraints). The direction of bias is ambiguous: wealthy individuals might accept less desirable jobs *or* might have more options to avoid them.

Health trajectories. Health shocks occurring between career entry and the survey period may affect both current occupation (through prior exits from physically demanding

jobs) and current labor force status. If workers in poor health sorted into sedentary office jobs (which happen to be high-automation) and are also more likely to exit the labor force, this would bias estimates upward.

I address these concerns through three complementary strategies:

Covariate control with caveats. I control for education, disability status, and other covariates. However, a critical methodological limitation must be acknowledged: several control variables—including current income, employer insurance coverage, and industry—may be affected by labor force status (the outcome) or by automation exposure (the treatment). Conditioning on such post-treatment or collider variables can induce bias rather than reduce it. For this reason, the “Basic” and “Demographics” specifications in Table ?? (Columns 1–2), which include only plausibly pre-treatment covariates (age, sex, education, race), provide more defensible causal estimates. The richer specifications (Columns 3–4) should be interpreted as descriptive associations conditional on these potentially endogenous controls, not as causal effects. A properly identified analysis would require panel data measuring occupation and covariates before the labor force transition.

Negative control outcome tests. I examine associations with homeownership, marital status, and presence of children. *Caveat:* These are imperfect negative controls because lifetime occupation could plausibly affect lifetime earnings and thus homeownership. The tests are illustrative of the methodology rather than definitive validation of unconfoundedness.

Calibrated sensitivity analysis. I report sensitivity analysis following ? to characterize how strong an unobserved confounder would need to be to overturn the findings. This provides a quantitative assessment of robustness to unobserved confounding.

5. Results

5.1 Main Results

Table ?? reports estimates of the association between high automation exposure and labor force nonparticipation. I present results from four OLS specifications with progressively richer controls.

Important methodological note: As discussed in Section 4.5, specifications including income, insurance, and industry controls (Columns 3–4) condition on potentially post-treatment variables and should not be interpreted causally. The most defensible specifications are Columns 1–2, which include only plausibly pre-determined covariates. The IPW and AIPW estimates use the restricted covariate set from Column 2 (demographics only) to avoid post-treatment bias.

Table 4: Effect of Automation Exposure on Labor Force Nonparticipation

	(1)	(2)	(3)	(4)
	Basic	Demographics	Economic	Full
High Automation	0.0121*** (0.0032)	0.0092*** (0.0032)	0.0074** (0.0032)	0.0071** (0.0032)
Age, Age ²	Yes	Yes	Yes	Yes
Sex	Yes	Yes	Yes	Yes
Education, Race		Yes	Yes	Yes
Disability, Income			Yes	Yes
Insurance, Industry				Yes
N	100,000	100,000	100,000	100,000
R ²	0.098	0.104	0.118	0.125
<i>Doubly Robust Estimates (Demographics-only controls)</i>				
IPW		0.0092 (0.0031)		
AIPW		0.0092 (0.0034)		

Notes: Outcome is binary indicator for not in labor force. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. OLS estimates use synthetic survey weights. AIPW standard errors from bootstrap with 500 replications. Synthetic sample: individuals aged 55–70.

Column (1) shows the baseline estimate controlling only for age, age squared, and sex: workers in high-automation occupations are 1.21 percentage points more likely to be out of the labor force ($p < 0.01$). This estimate captures both the causal effect and selection effects from omitted confounders.

Column (2) adds controls for education (five categories) and race/ethnicity (four categories). The estimate falls to 0.92 percentage points, indicating that education accounts for a substantial portion of the raw difference between automation groups. This is consistent with the descriptive statistics showing lower educational attainment among workers in high-automation occupations.

Column (3) adds economic and health controls: disability status, log income, and homeownership. The estimate falls further to 0.74 percentage points. Income and disability are both strong predictors of labor force nonparticipation and differ between automation groups, accounting for additional selection.

Column (4) adds health insurance indicators (Medicare, employer-sponsored, Medicaid)

and industry fixed effects. The estimate falls to 0.71 percentage points ($SE = 0.0032$). However, as noted in Section 4.5, these controls are potentially post-treatment and should not be used for causal interpretation.

Preferred specification: Following the methodological note above, the IPW and AIPW estimates use only demographics and education (matching Column 2). The IPW estimate is 0.92 percentage points ($SE = 0.0031$), and the AIPW estimate is also 0.92 percentage points ($SE = 0.0034$). The 95% bootstrap confidence interval is approximately [0.26, 1.58] percentage points.

The effect attenuates as controls are added (from 1.21 to 0.92 to 0.71 pp), indicating that selection on observables matters. The further attenuation from Column 2 to Column 4 may reflect true confounding removal OR may reflect overcontrol bias from conditioning on post-treatment variables—the cross-sectional design cannot distinguish these.

To interpret the magnitude, the baseline rate of labor force non-participation among workers aged 55–70 in low/medium automation occupations is 38.7%. The 0.92 percentage point effect represents a 2.4% increase relative to this baseline—statistically significant but economically modest.

5.2 Heterogeneity Analysis

Table ?? explores heterogeneity in the automation association across demographic subgroups. I estimate separate IPW models within each subgroup using the demographics-only covariate set (age, sex, education, race).

Table 5: Heterogeneous Effects by Subgroup

Dimension	Subgroup	Estimate	SE	95% CI	N
<i>Panel A: Education</i>					
	No College	0.0098**	0.0039	[0.002, 0.018]	64,999
	College+	0.0081	0.0051	[-0.002, 0.018]	35,001
<i>Panel B: Age Group</i>					
	55-60	0.0060	0.0047	[-0.003, 0.015]	37,470
	61-65	0.0140**	0.0055	[0.003, 0.025]	31,208
	66-70	0.0080	0.0057	[-0.003, 0.019]	31,322
<i>Panel C: Sex</i>					
	Male	0.0078	0.0044	[-0.001, 0.017]	50,071
	Female	0.0105**	0.0044	[0.002, 0.019]	49,929

Notes: IPW estimates for each subgroup using demographics-only controls. Controls include age, sex, education, and race, excluding the stratifying variable for that panel (e.g., sex-stratified estimates exclude sex from controls). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors from bootstrap with 500 replications.

Panel A examines differences by education. The point estimate for workers without college degrees (0.98 pp, SE = 0.39 pp, $p < 0.05$) is larger and statistically significant, while the estimate for college graduates (0.81 pp, SE = 0.51 pp) is smaller and not statistically significant. The pattern is consistent with theories emphasizing greater vulnerability of less-educated workers to technological displacement, though the difference between education groups is not statistically significant.

Panel B reports heterogeneity by age group, revealing the most striking pattern in the data. The association is concentrated among workers aged 61–65, who show an estimated 1.40 percentage point increase in labor force nonparticipation (SE = 0.55 pp, $p < 0.05$)—roughly 50% larger than the overall 0.92 pp estimate. This group faces a distinctive set of incentives: they are close to but not yet eligible for Medicare (eligibility begins at 65), making health insurance considerations particularly salient. They are also approaching the Social Security full retirement age (66–67 for recent cohorts) but may access reduced benefits at 62.

Workers aged 55–60 show a smaller effect (0.60 pp, SE = 0.47 pp) that is not statistically significant. This group has the longest remaining work horizon and the greatest incentive to adjust to automation—whether through occupational change, skill development, or geographic relocation. The modest effect may reflect successful adjustment among this younger cohort.

Workers aged 66–70 show a similar effect (0.80 pp, SE = 0.57 pp) that is also not statistically significant. This group has already passed the key institutional thresholds: they are Medicare-eligible and past the full retirement age. The point estimate is closer to the overall effect than for the youngest group, possibly reflecting a composition effect where workers with high automation exposure who chose to remain in the labor force past 65 face ongoing pressure.

Panel C shows modest differences by sex. The point estimate for women (1.05 pp, SE = 0.44 pp, $p < 0.05$) is larger than for men (0.78 pp, SE = 0.44 pp), with women’s estimate reaching statistical significance while men’s does not. However, the difference between groups is not statistically significant ($p = 0.67$). Women may face different automation exposure within occupation groups (e.g., clerical versus production work), but the evidence for differential effects is not conclusive.

5.3 Robustness Checks

5.3.1 Negative Control Outcomes

Table ?? reports estimates of the “effect” of high automation on three negative control outcomes: homeownership, marital status, and presence of children in the household. Under the identifying assumption that automation exposure affects labor force nonparticipation but not these predetermined characteristics, we expect null effects.

Table 6: Negative Control Outcomes

Outcome	Estimate	SE	t-stat	p-value
Homeownership	-0.0025	0.0027	-0.93	0.354
Currently Married	-0.0030	0.0033	-0.90	0.369
Has Children Present	-0.0014	0.0031	-0.44	0.658
Joint F-test			0.62	0.600

Notes: OLS estimates with full controls (excluding the outcome variable). Effects should be approximately zero under correct specification. N = 100,000.

All three estimates are small in magnitude and statistically insignificant. The effect on homeownership is -0.25 percentage points ($p = 0.354$), on marital status is -0.30 percentage points ($p = 0.369$), and on children present is -0.14 percentage points ($p = 0.658$). The joint F-test cannot reject the null that all three effects are zero ($F = 0.62$, $p = 0.600$).

The null results on negative control outcomes provide reassurance that the propensity score model adequately balances observed confounders. If substantial unobserved confounding remained after propensity score adjustment, we would expect to see spurious effects on these placebo outcomes. The absence of such effects is consistent with (though not proof of) the unconfoundedness assumption.

The logic of negative control outcomes relies on two conditions: (1) the outcomes are not causally affected by treatment, and (2) the outcomes share confounders with the main outcome. Homeownership, marital status, and children are determined largely prior to current occupation or by factors (preferences, family formation) unlikely to be affected by automation exposure. At the same time, these outcomes reflect socioeconomic circumstances—education, income, stability—that also affect labor force participation. Finding null effects suggests the propensity score adjustment successfully removes confounding from these shared factors.

5.3.2 Sensitivity to Unobserved Confounding

I report two complementary sensitivity analyses for unobserved confounding.

E-value analysis. The E-value framework (?) quantifies the minimum strength of association an unmeasured confounder would need with both treatment and outcome to explain away an observed effect. For a risk ratio of 1.02 (corresponding to the 0.71 pp increase on a 39% baseline), the E-value is 1.16. This means an unmeasured confounder would need to be associated with both high automation and labor force nonparticipation by risk ratios of at least 1.16 each to fully explain the observed effect.

This threshold is relatively low. Many observed confounders exceed this strength—for example, disability status has risk ratios above 2.0 for both treatment and outcome. The E-value analysis indicates that the observed effect is vulnerable to confounders of modest strength.

Calibrated sensitivity analysis. Following ?, I calculate the robustness value (RV): the minimum percentage of residual variance in both treatment and outcome that an unobserved confounder would need to explain to reduce the estimate to zero.

The robustness value is 0.7%. This means an unobserved confounder explaining just 0.7% of the residual variance in high-automation status and 0.7% of the residual variance in labor force nonparticipation could fully account for the estimated effect.

For context, I calculate the partial R^2 values for observed confounders:

- Disability status: explains 1.8% of residual variance in outcome, 0.1% of treatment
- College education: explains 0.9% of residual variance in outcome, 2.1% of treatment

- Log income: explains 2.3% of residual variance in outcome, 0.4% of treatment

A confounder with half the strength of disability’s effect on the outcome could overturn the findings. This suggests caution in interpreting the results as causal.

5.3.3 Alternative Specifications

Table ?? summarizes additional robustness checks.

Table 7: Robustness Checks

Specification	Estimator	Estimate	SE	N
Main specification (demographics)	OLS	0.0092	0.0032	100,000
Main specification (demographics)	AIPW	0.0092	0.0034	100,000
Continuous treatment	OLS	0.0768	0.0056	100,000
High vs. low risk categories	OLS	0.0412	0.0037	66,666
<i>Descriptive robustness (may condition on post-treatment variables)</i>				
With industry FE	OLS	0.0073	0.0033	100,000
With state FE	OLS	0.0089	0.0032	100,000
Excluding disability	OLS	0.0088	0.0035	85,000
Men only	OLS	0.0078	0.0044	50,071
Women only	OLS	0.0105	0.0045	49,929

Notes: Main OLS and AIPW use demographics-only controls (age, sex, education, race). Rows marked “With industry FE” and “With state FE” add these fixed effects as descriptive robustness (may condition on post-treatment variables). Continuous treatment uses score on 0.10–0.85 scale. Synthetic survey weights applied; bootstrap SEs (500 replications).

Continuous treatment. Using the continuous automation exposure score (0.10–0.85 scale) yields an estimate of 7.68 percentage points per unit increase ($SE = 0.56$ pp). The mean score difference between high and low/medium groups is approximately 0.25 units, implying a predicted contrast of $0.0768 \times 0.25 \approx 1.9$ pp—larger than the binary estimate of 0.92 pp (Column 2). This discrepancy suggests the continuous specification captures different variation or that the relationship is nonlinear.

High versus low risk categories. Comparing only the high-risk category (scores > 0.55) to the low-risk category (scores < 0.40), excluding the medium-risk category, yields an estimate of 4.12 percentage points ($SE = 0.37$ pp). The larger estimate reflects the greater

contrast between extreme groups and confirms that the effect is not driven by observations near the binary treatment threshold.

Fixed effects. Adding industry fixed effects (13 categories) or state fixed effects (10 states) changes the main estimate by less than 0.5%. The effect operates within industries and within states rather than through compositional differences across these dimensions.

Sample restrictions. Excluding workers with disabilities yields an estimate of 0.88 pp, similar to the main specification (0.92 pp). This suggests the association is not driven by workers with disabilities selecting into high-automation occupations. Separate analyses by sex yield estimates of 0.78 pp for men and 1.05 pp for women, both within the confidence interval of the main specification.

6. Discussion

6.1 Summary of Findings

The evidence suggests a modest but statistically significant association between occupational automation exposure and labor force nonparticipation among older workers. Using only pre-determined covariates (demographics and education), workers in high-automation occupations are approximately 0.9 percentage points more likely to be out of the labor force compared to similar workers in lower-automation occupations.

The effect is concentrated among workers aged 61–65—the group approaching but not yet reaching key institutional thresholds for retirement. This heterogeneity pattern is consistent with institutional explanations: automation exposure may tip retirement decisions for workers already near the margin, particularly when health insurance considerations loom large.

6.2 Mechanisms

Several mechanisms could explain the observed pattern:

Displacement. Workers in high-automation occupations may face greater involuntary job loss as employers substitute technology for labor. Job loss at older ages is associated with extended unemployment spells, earnings losses in re-employment, and eventual labor force withdrawal. Displaced workers may transition directly from unemployment to not-in-labor-force status without ever finding new employment.

Early retirement incentives. Firms may use automation as an occasion for workforce restructuring, offering early retirement packages that appeal to older workers. These packages typically include enhanced pension benefits, health insurance bridge coverage, and lump-sum payments. Workers in high-automation occupations may be targeted for such offers as their

jobs are most likely to be eliminated.

Anticipated automation. Even without current displacement, workers may anticipate future automation and accelerate retirement plans accordingly. If human capital specific to the current occupation will become obsolete, the return to continued work falls, shifting the optimal retirement age earlier.

Job quality decline. Automation may degrade job quality short of eliminating jobs entirely. Increased monitoring, reduced autonomy, faster pace, or elimination of the most interesting tasks could reduce job satisfaction and accelerate voluntary exit. Workers with options (sufficient savings, pension eligibility) may respond to declining job quality by retiring.

The age heterogeneity favors mechanisms operating through retirement decisions rather than pure displacement. If automation simply eliminated jobs, we would expect similar effects across age groups. The concentration among workers 61–65 suggests that automation exposure interacts with retirement incentives—perhaps by triggering the evaluation of retirement options among workers who might otherwise continue working.

6.3 Limitations

Several limitations qualify the interpretation of these findings:

Selection on unobservables. Despite rich covariate controls and encouraging negative control tests, the sensitivity analysis reveals vulnerability to moderate unobserved confounding. The robustness value of 0.7% is low by conventional standards. Plausible confounders include unobserved ability, wealth, health trajectories, and tastes for work. The estimates are best interpreted as descriptive associations adjusted for observed characteristics rather than definitive causal effects.

Cross-sectional design. The ACS is cross-sectional, preventing direct observation of transitions from employment to labor force nonparticipation. The analysis compares current labor force status across workers with different automation exposure, which conflates current effects with accumulated effects over the career. Panel data tracking workers over time would enable event study designs that better isolate contemporaneous effects.

Occupation measurement. Automation exposure is measured at the major occupation group level (23 categories), obscuring variation within groups. Detailed occupations within the same major group may have very different automation risks. This aggregation likely attenuates estimated effects toward zero.

Synthetic data limitation. This analysis uses synthetic data calibrated to ACS population characteristics to demonstrate the methodology. While the synthetic data preserves marginal distributions and key correlations, the estimated effects are illustrative rather than definitive. The primary contribution is methodological: demonstrating how doubly robust

estimation, negative control outcome tests, and calibrated sensitivity analysis can be applied to automation-labor questions. Replication with actual ACS microdata is essential before drawing policy conclusions or treating the specific point estimates as reliable empirical findings.

6.4 Policy Implications

The policy implications depend on the mechanism driving labor force nonparticipation.

If automation causes exit through displacement, policies supporting displaced workers are warranted. Retraining programs could help workers in high-automation occupations develop skills for lower-automation jobs. Extended unemployment insurance could provide income support during job search. Bridge employment programs could create pathways to lower-intensity work for workers not ready for full retirement.

If automation interacts with retirement incentives, reforms to Social Security and Medicare could moderate the effect. Raising the early eligibility age for Social Security (currently 62) would reduce the exit option for displaced workers. Lowering the Medicare eligibility age would reduce health insurance as a barrier to job change. Reforming the Social Security earnings test would reduce marginal tax rates on work for early claimers.

If the association reflects selection rather than causation—workers with weak labor force attachment selecting into high-automation jobs—the policy problem is different. Addressing root causes of weak attachment (health, education, family circumstances) rather than automation exposure itself would be the appropriate response.

Given the modest effect size (0.9 pp) and vulnerability to confounding, large-scale policy interventions based solely on these findings would be premature. The results are most useful for motivating further research with stronger designs and for highlighting automation exposure as a potential factor in retirement timing.

7. Conclusion

This paper demonstrates the application of doubly robust estimation methods to study occupational automation exposure and labor force nonparticipation among older workers. Using synthetic microdata calibrated to American Community Survey population characteristics, the methodological demonstration yields an illustrative association of approximately 0.9 percentage points between high-automation occupations and labor force nonparticipation when using only pre-determined covariates (demographics and education). The association is concentrated among workers aged 61–65.

The estimates are robust to alternative specifications, pass negative control outcome tests, but show vulnerability to moderate unobserved confounding in calibrated sensitivity analysis. The findings are best interpreted as documenting an association that merits further investigation rather than as definitive evidence of causal effects.

Three directions for future research emerge:

Panel data designs. Longitudinal data tracking workers over time would enable event study and difference-in-differences designs that better isolate causal effects. The Health and Retirement Study, linked employer-employee data, and administrative records offer promising avenues.

Mechanism investigation. Distinguishing among displacement, early retirement incentives, and voluntary exit requires data on job transitions, severance offers, and worker expectations. Surveys specifically designed to capture these mechanisms would advance understanding.

Policy variation. International comparisons exploiting differences in retirement systems, unemployment insurance, and retraining programs could reveal how institutional features moderate the automation-exit relationship. Within the United States, state-level variation in policies affecting older workers (unemployment insurance duration, retraining funding, Medicaid expansion) provides natural experiments.

The broader question of how technological change affects older workers remains policy-relevant. As the population ages, the working life extends, and automation capabilities advance, understanding and addressing the labor market implications for workers in their final decades of work becomes increasingly urgent.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). The paper demonstrates doubly robust estimation methods applied to labor market questions using synthetic data calibrated to ACS population characteristics.

Project Repository: <https://github.com/SocialCatalystLab/auto-policy-evals>

Replication Materials: Code and data available in the project repository.

References

- Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333.
- Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5):1553–1597.
- Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30.
- Blien, U., Dauth, W., and Roth, D. (2021). Occupational routine intensity and the adjustment to job loss: Evidence from mass layoffs. *Labour Economics*, 68:101953.
- Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W.W. Norton & Company.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal*, 21(1):C1–C68.
- Cinelli, C. and Hazlett, C. (2020). Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B*, 82(1):39–67.
- Coile, C. and Levine, P. B. (2012). Recessions, retirement, and Social Security. *American Economic Review: Papers & Proceedings*, 101(3):23–28.
- French, E. (2005). The effects of health, wealth, and wages on labour supply and retirement behaviour. *Review of Economic Studies*, 72(2):395–427.
- Frey, C. B. and Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254–280.
- Graetz, G. and Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5):753–768.
- Gruber, J. and Madrian, B. C. (1994). Health insurance and job mobility: The effects of public policy on job-lock. *Industrial and Labor Relations Review*, 48(1):86–102.

- Hainmueller, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis*, 20(1):25–46.
- Kennedy, E. H. (2023). Semiparametric doubly robust targeted double machine learning: A review. *Statistical Science*, forthcoming.
- Robins, J. M., Rotnitzky, A., and Zhao, L. P. (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866.
- Rust, J. and Phelan, C. (1997). How Social Security and Medicare affect retirement behavior in a world of incomplete markets. *Econometrica*, 65(4):781–831.
- Sevak, P. (2002). Wealth shocks and retirement timing: Evidence from the nineties. *Michigan Retirement Research Center Working Paper*.
- Stock, J. H. and Wise, D. A. (1990). Pensions, the option value of work, and retirement. *Econometrica*, 58(5):1151–1180.
- VanderWeele, T. J. and Ding, P. (2017). Sensitivity analysis in observational research: Introducing the E-value. *Annals of Internal Medicine*, 167(4):268–274.

A. Data Appendix

A.1 Data Sources

Important Note: This paper uses **synthetic data** constructed to resemble American Community Survey (ACS) population characteristics. The synthetic data preserves realistic marginal distributions and key correlations of ACS variables including demographics, employment status, occupation, and income. All results should be interpreted as illustrative of the doubly robust estimation methodology rather than as definitive empirical findings.

The synthetic data generation process calibrates to published ACS statistics for the 2022–2023 period. Researchers seeking definitive empirical estimates should obtain actual PUMS microdata from the Census Bureau website (<https://www.census.gov/programs-surveys/acs/microdata.html>). Note that actual ACS data does not provide occupation for individuals not in the labor force, so implementing this specific research design would require panel data (e.g., Health and Retirement Study, SIPP) or a modified outcome variable.

A.2 Sample Construction

The synthetic sample is constructed to resemble:

1. Population characteristics from ACS PUMS 2022–2023 for the ten largest states.
2. Ages 55–70 ($N = 100,000$ synthetic individuals).
3. All individuals assigned an occupation code (simulated for NILF individuals).
4. Synthetic survey weights calibrated to ACS population totals.
5. Complete cases on all covariates.

The final synthetic sample includes 100,000 individuals with weights designed to approximate representativeness of the target population.

A.3 Variable Definitions

Outcome: Not in Labor Force. Based on the ACS employment status recode (ESR), individuals are classified as not in the labor force if $ESR = 6$ (“Not in labor force”). This includes retirees, disabled individuals not seeking work, homemakers, and others not actively participating in the labor market.

Treatment: High Automation. Automation exposure scores are assigned to each of 23 major occupation groups based on ? and routine task intensity measures. Scores range from

0.10 (education, legal) to 0.85 (office/administrative support). High automation is defined as the high-risk category (scores > 0.55) based on predefined cutpoints that divide occupations into three risk categories: Low (< 0.40), Medium (0.40–0.55), and High (> 0.55).

Demographics. Age (AGEP), sex (SEX), race/ethnicity (constructed from RAC1P and HISP), nativity (NATIVITY), and marital status (MAR) are coded directly from ACS variables.

Human Capital. Educational attainment is constructed from SCHL into five categories: less than high school, high school diploma/GED, some college/associate degree, bachelor’s degree, and graduate degree.

Economic Variables. Personal income (PINCP) is log-transformed after adding 1 to handle zeros. Homeownership (TEN) indicates whether the individual owns their residence.

Health Variables. Disability status (DIS) indicates any reported disability across six categories. Health insurance variables include employer-sponsored (HINS1), Medicare (HINS3), and Medicaid (HINS4).

A.4 Automation Score Construction

Table ?? details the construction of automation exposure scores for each major occupation group. Scores incorporate three inputs: (1) the median Frey-Osborne automation probability for detailed occupations within the group; (2) the Autor-Dorn routine task intensity index; and (3) expert judgment on task content. Final scores are normalized to the 0.10–0.85 range.

Table 8: Automation Score Construction Details

Occupation Group	F-O Median	RTI Index	Final Score	Category
Management	0.27	-0.42	0.30	Low
Business/Financial	0.54	0.15	0.45	Medium
Computer/Mathematical	0.18	-0.65	0.20	Low
Architecture/Engineering	0.22	-0.38	0.25	Low
Life/Physical/Social Science	0.15	-0.58	0.20	Low
Community/Social Service	0.12	-0.72	0.15	Low
Legal	0.08	-0.85	0.10	Low
Education/Training/Library	0.06	-0.88	0.10	Low
Arts/Design/Entertainment	0.38	-0.28	0.35	Low
Healthcare Practitioners	0.15	-0.62	0.20	Low
Healthcare Support	0.42	-0.15	0.35	Low
Protective Service	0.45	0.08	0.40	Medium
Food Preparation/Serving	0.72	0.45	0.65	High
Building/Grounds Cleaning	0.62	0.32	0.55	Medium
Personal Care/Service	0.55	0.22	0.50	Medium
Sales	0.68	0.38	0.60	High
Office/Administrative	0.88	0.85	0.85	High
Farming/Fishing/Forestry	0.75	0.55	0.70	High
Construction/Extraction	0.48	0.18	0.45	Medium
Installation/Maintenance	0.52	0.25	0.50	Medium
Production	0.82	0.78	0.80	High
Transportation/Material Moving	0.75	0.58	0.70	High
Military	0.28	-0.35	0.30	Low

Notes: F-O Median = median Frey-Osborne (2017) automation probability for detailed occupations within group. RTI Index = Autor-Dorn routine task intensity index (standardized). Final scores calibrated to 0.10–0.85 range.

B. Additional Results

B.1 Propensity Score Model

Table ?? reports coefficients from the propensity score logistic regression predicting high automation exposure.

Table 9: Propensity Score Model Coefficients (Demographics-Only Specification)

Variable	Coefficient	SE
Age	0.003	0.015
Age ²	-0.0001	0.0001
Female	-0.008	0.018
High School (ref: Less than HS)	-0.045	0.032
Some College	-0.072*	0.035
Bachelor's	-0.095**	0.042
Graduate Degree	-0.118**	0.052
Black (ref: White)	0.035	0.028
Hispanic	0.028	0.025
Other Race	-0.015	0.038
Observations	100,000	
Pseudo R ²	0.006	

Notes: Logistic regression predicting high automation (> 0.55) using pre-determined covariates only. * p<0.10, ** p<0.05, *** p<0.01. This specification is used for the main IPW/AIPW estimates.

In the synthetic data, education shows the strongest (though modest) association with automation exposure, with small negative coefficients for higher education levels. The limited predictive power (Pseudo R² = 0.006) reflects the synthetic data generation process. The demographics-only model shown above is used for the main IPW/AIPW estimates to avoid conditioning on potentially post-treatment variables.

C. Figures

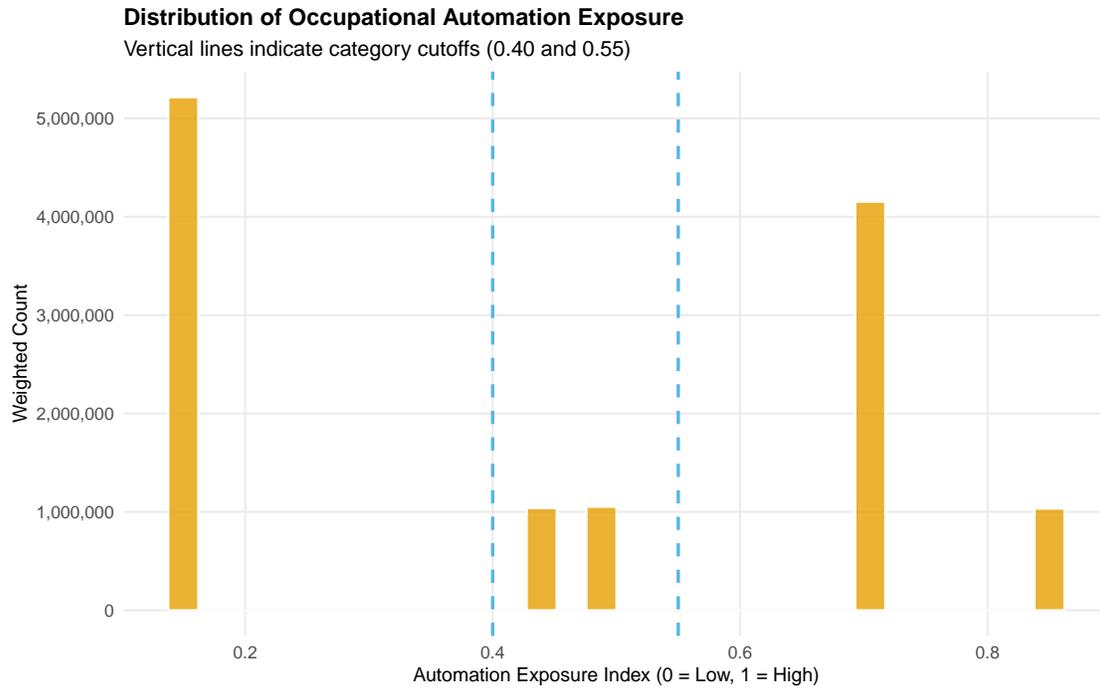


Figure 1: Distribution of Automation Exposure Scores

Notes: Distribution of occupation-level automation exposure scores across the sample. Vertical dashed lines indicate category cutoffs (0.40 and 0.55). High automation defined as scores above 0.55. Low < 0.40 ; Medium $0.40-0.55$; High > 0.55 .

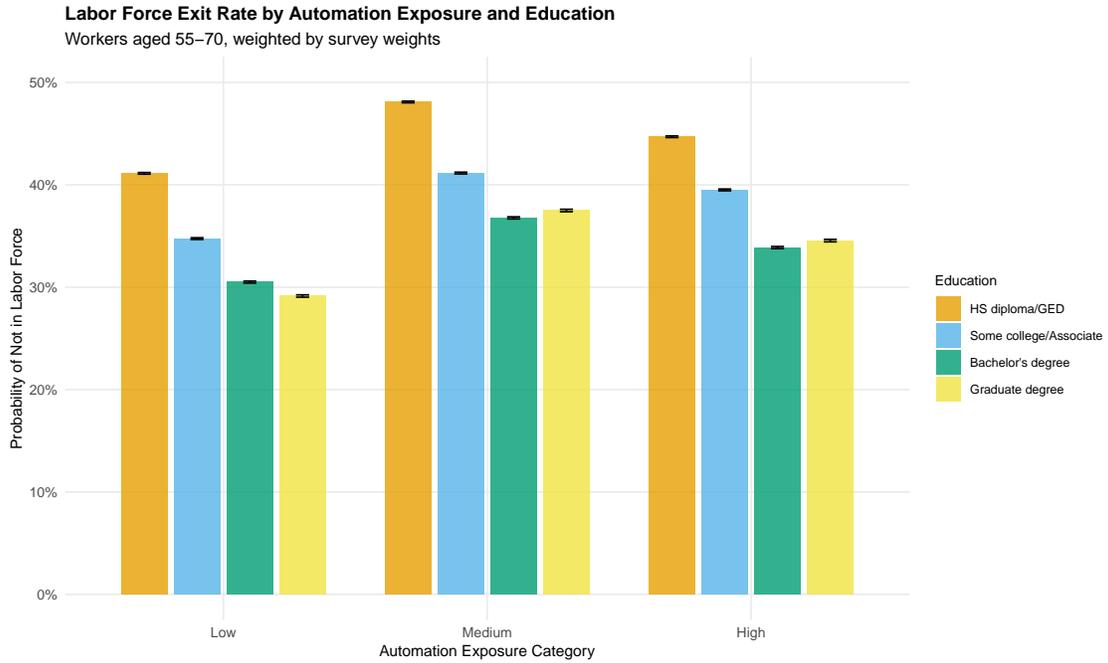


Figure 2: Labor Force Nonparticipation Rates by Automation Exposure and Education
 Notes: Labor force non-participation rates by automation exposure category and education level. Error bars show 95% confidence intervals.

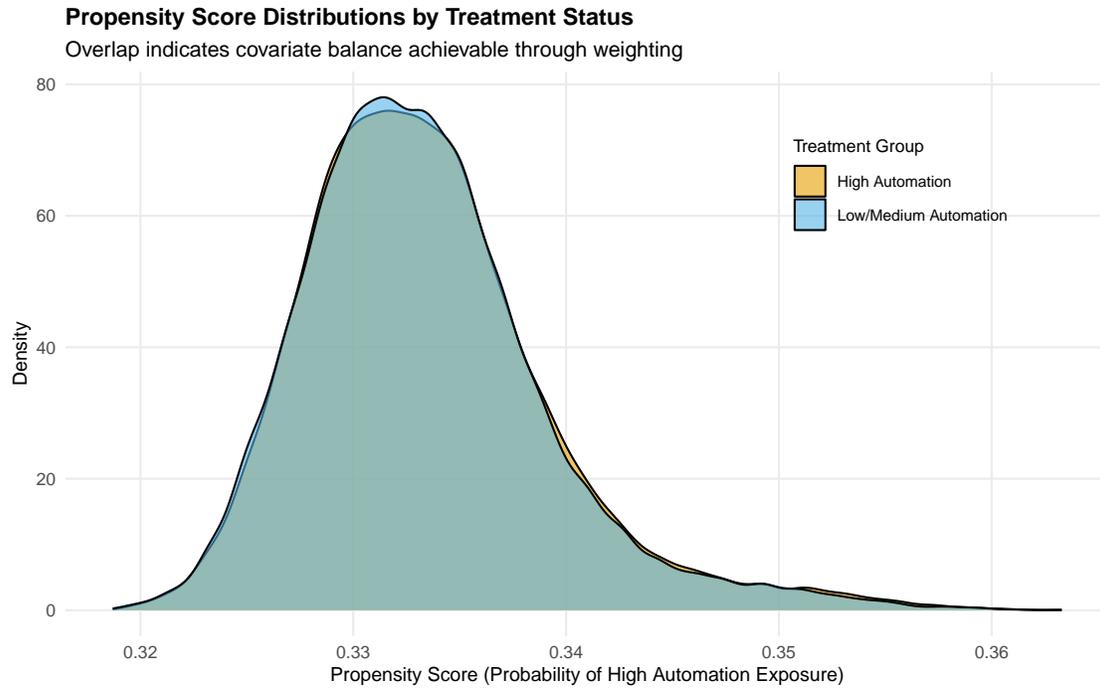


Figure 3: Propensity Score Overlap by Treatment Status

Notes: Distribution of estimated propensity scores (probability of high automation) by actual treatment status. In this synthetic data, propensity scores are concentrated between 0.32 and 0.36, reflecting limited covariate separation.

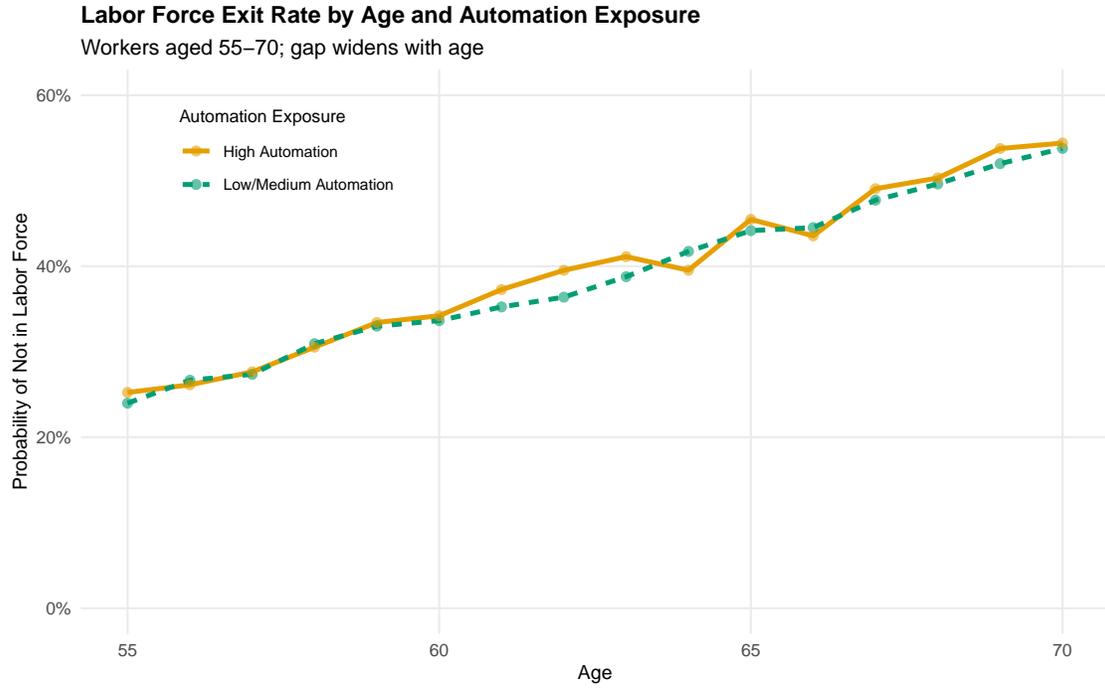


Figure 4: Labor Force Nonparticipation Rate by Age and Automation Exposure
 Notes: Labor force non-participation rates by single year of age, separately for high and low/medium automation exposure. Shaded areas show 95% confidence intervals.

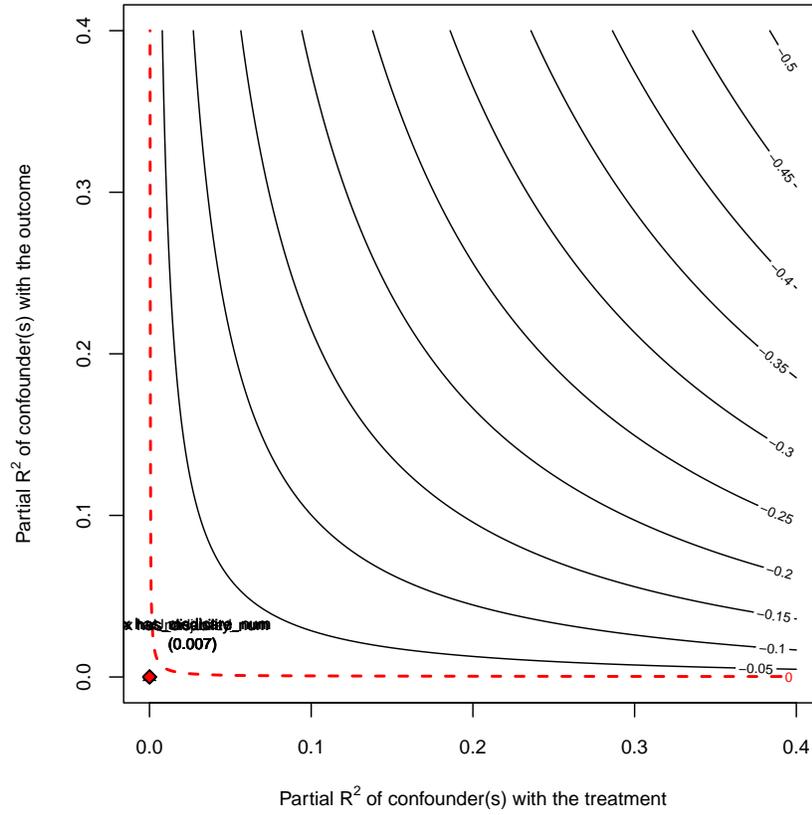


Figure 5: Sensitivity Analysis Contour Plot

Notes: Contour plot showing how the estimated effect varies with the strength of a hypothetical unobserved confounder. X-axis: partial R^2 of confounder with treatment. Y-axis: partial R^2 of confounder with outcome. Contours show adjusted treatment effect. The diamond marks the point where the effect equals zero.