

参赛学生姓名：陈景羲，胡越

中学：上海中学国际部

省份：上海

国家/地区：中国

指导老师姓名：何天舟

指导老师单位：上海中学国际部

论文题目：高薪与慢增长：职业 AI 暴露与中国中年收入风险——来自 CFPS 2020–2022 的证据

Higher Pay, Slower Growth: Occupational AI Exposure and Mid-Career Earnings Risk in China — Evidence from CFPS 2020–2022

Abstract

In the modern economy, occupational AI exposure indices like those from the OECD (Organization for Economic Co-operation and Development) categorize labor at the job-title level, yet the same title can obscure vastly different earnings stakes for mid-career college graduates and less-educated workers in China. By attaching the 2020 and 2022 waves of the China Family Panel Studies to year-specific occupational scores (Zhang et al.), the Ministry of Science and Technology rollout of National New Generation AI Innovation and Development Pilot Zones, and the 2019 provincial totals of the Peking University Digital Financial Inclusion Index and AI patents, we disentangle individual-level vulnerability from external factors such as industrial policy.

First, demographically, each additional schooling level increases exposure by about 0.048, with college-educated workers aged 35–49 carrying a further premium of 0.062. Second, in terms of pay trajectory, higher exposure correlates with 0.64 higher log income after demographic and occupation-group controls. However, baseline 2020 exposure predicts about 0.23 to 0.25 lower Δ log income through 2022, reflecting task-based substitution within cognitive-heavy occupations during early AI capability expansion. In other words, this illustrates the “Higher Pay, Slower Growth” paradox. Third, in terms of geographic policy shocks, our results show that newly designated pilot provinces exhibit only a statistically small average growth effect, while the exposure gradient remains large in the sample of clean difference-in-difference estimates (-0.30 to -0.38) and persists at -0.238 ($p < 0.01$) under predetermined provincial finance and patent controls. Consequently, individual earnings risk during this early expansionary AI capability window follows occupational task vulnerability rather than place-based policy implementation. Accordingly, policy interventions such as wage insurance, mid-career retraining, and career counseling should directly target individuals by demographic characteristics—schooling, mid-career status, and occupational exposure.

Keywords: AI exposure; staggered difference-in-differences; earnings risk; occupational tasks; China Family Panel Studies

Contents

1	Introduction	4
2	Literature Review	5
2.1	Occupational Exposure and Worker Heterogeneity	5
2.2	Skill-Biased Technical Change and Earnings in Levels versus Growth	6
2.3	AI Pilot Zones as Local Industrial Policy	6
2.4	Hypotheses	7
3	Data and Empirical Strategy	7
3.1	Data and Sample	7
3.2	Measuring AI Exposure	8
3.3	Policy Treatment: AI Pilot Zones	9
3.4	Empirical Strategy	9
4	Results	11
4.1	Who Holds High Exposure Scores	11
4.2	Pay in Levels, Slower Growth	12
4.3	Occupational Exposure and Place-Based Policy	16
4.4	Summary of Findings	20
5	Discussion and Policy	20
5.1	Policy Implications	20
5.2	Limitations	21
6	Conclusion	21
A	Additional Results: Nonlinear Predictors of AI Exposure	22
A.1	Models and Fit	22
A.2	Age and Education	23
	Acknowledgment	26

1 Introduction

Occupational artificial intelligence (AI) exposure scores rank jobs by how closely their tasks overlap with the capability of current artificial intelligence systems[1, 2, 3, 4]. In China, however, the same occupation title can encompass both a mid-career college graduate in a metropolitan area and a less-educated worker with dramatically lower household earnings. Because a score assigned strictly at the occupational level does not vary within job titles, standard job rankings fail to reveal earnings-risk susceptibility. Over the same period, the Ministry of Science and Technology (MOST) initiated a city-by-city rollout of National New Generation AI Innovation and Development Pilot Zones. Consequently, both place-based industrial policy and occupational task content can move worker earnings.

To disentangle those mechanisms, this paper structures its empirical investigation around three central questions:

1. **Demographic Allocation.** Once occupational AI exposure scores are attached to individuals in the China Family Panel Studies (CFPS), how strongly does exposure concentrate among better-educated workers in mid-career, and how does it line up with income in a cross section?
2. **“Higher Pay, Slower Growth” Phenomenon.** How does occupational exposure align with log income in cross sections, and how does higher 2020 exposure predict a slower change in log income ($\Delta \log \text{income}$) through 2022?
3. **“Task Vulnerability vs. Place-Based Policy.”** Is the observed earnings-growth penalty driven primarily by occupational task vulnerability, or are newly designated pilot provinces and regional digital developments the principal factors?

This paper contributes to the literature by utilizing a worker-side panel approach. Standard occupational indices abruptly stop at the job title [1, 2]. Existing studies of exposure, vacancies, and complementary skills rarely take a Chinese household panel into an income-growth design [3, 5, 6]. Meanwhile, pilot zone evaluations focus almost exclusively on firm-level outcomes within designated cities, rather than tracking individual worker earning trajectories [7, 8]. Ultimately, we address the deficiencies of these literatures by linking the 2020 and 2022 waves of the CFPS to year-specific occupational exposure scores from Zhang et al.[9, 4]. Pivotaly, our fieldwork was largely completed before the mass dissemination of large language models like ChatGPT. Thus, our model captures an early generative-AI capability window rather than a post-adoption equilibrium.

- **First-Difference Estimates.** By regressing log income changes ($\Delta \log \text{income}$) on 2020 occupational exposure, we remove time-invariant intercepts at the individual level.
- **Difference-in-Differences (DID).** Dropping always-treated provinces (designated by 2020) isolates a clean comparison. Parallel comparisons are conducted at the city level.
- **Leave-out Occupation Bartik and Predetermined Controls (Robustness).** We instrument individual exposure with the leave-out mean exposure in the same two-digit ISCO cell. We then add the 2019 provincial Peking University Digital Financial Inclusion Index (DFIIC) and 2019 provincial AI patent counts as controls for predetermined digital financial development and patent intensity.

To summarize our findings:

1. **Allocation in Levels.** Education raises exposure by about 0.048 for each additional category on the survey scale. College-educated workers aged 35–49 have a further premium of about 0.062. Group means are near 0.60 versus 0.40 for less-educated workers outside that age band.

2. **Pay in Levels vs. Slower Growth.** Then, among those with positive income, the same exposure still correlates 0.64 higher log income after demographics and occupation-group indicators. The growth side is different. Baseline exposure predicts about 0.23 to 0.25 lower Δ log income over 2020–2022. Putting the 2018 score on the same 2020 occupation leaves a similar penalty; a 2018–2020 growth regression does not.
3. **Task Exposure vs. Place-Based Policy.** Newly designated pilot provinces show only a small average growth effect. In the sample excluding always-treated provinces, the exposure coefficient is about -0.30 to -0.38 . After adding the 2019 provincial controls, it remains about -0.24 ($p < 0.01$), suggesting that the growth penalty is more closely associated with occupational task exposure than with place-based policy.

Collectively, these findings illustrate that individual earnings risk during this early generative-AI window follows task vulnerability rather than regional pilot labels or digital readiness. The remainder of the paper details our methodology to attain such a result and is organized as follows: Section 2 reviews the measurement, automation, and pilot-zone literatures and states three hypotheses. Data and designs follow in Section 3, empirical results across demographic allocation, levels versus growth, and place-based policy checks are completed in Section 4, and policy implications and limitations are addressed in Section 5.

2 Literature Review

The literature on AI exposure has moved from computerization forecasts to task-based occupational indices [1, 2, 10, 4]. A second line of labor economics examines how automation and AI reshape labor demand and wages [11, 12, 13, 14]. In China, a third empirical literature field now evaluates the National New Generation AI Innovation and Development Pilot Zones to study firm innovation and related outcomes after staggered designation [7, 8]. This review synthesizes three paradigms to ground our worker panel investigation across three empirical dimensions.

2.1 Occupational Exposure and Worker Heterogeneity

Whereas early automation studies primarily evaluated job categories to produce occupational-level rankings, subsequent frameworks have turned to generating binary forecasts with ability mappings [1]. Felten, Raj, and Seamans, in particular, matched progress on AI benchmarks with occupational abilities and then with standard occupational classification titles, and released occupation, industry, and county scores [2]. Tolan et al. score occupations against suites of cognitive AI tests rather than a single computerization probability [10]. For China, Zhang et al. built a year-specific SOC table from task ratings of the large language model and documented labor-demand responses at the occupation level [4]. These measurement tools consistently find high AI exposure in information-intensive, cognitive professions rather than routine manual work [3, 2].

Nevertheless, an occupation-level score assigned strictly at the job level does not reveal who occupies those jobs. Eloundou, Manning, Mishkin, and Rock estimate that a substantial share of U.S. employment is in jobs where large language models can take over a non-trivial fraction of tasks, and that higher-wage occupations are more exposed [3]. Acemoglu, Autor, Hazell, and Restrepo find that U.S. establishments whose task mix is closer to current AI capabilities posted more AI vacancies and fewer non-AI vacancies over 2010–2018. That said, occupation- and industry-level employment and wage effects were still minimal to detect [5]. Furthermore, Deming demonstrates that the labor market increasingly rewards social skills that machines cannot easily replace, meaning earnings paths need not move one-for-one with a title-level score [6].

The China Family Panel Studies follows the same people across waves and keeps occupation codes, schooling, age, and household income in one file [15, 9]. Attaching a year-specific score to those respondents lets us test how exposure concentrates among higher-educated workers in mid-career, relative to younger graduates or

less-educated workers in the same titles. Group differences in AI use and adaptation have been mapped. Whether baseline individual exposure predicts wage growth within a nationally representative household remains an open question, though.

2.2 Skill-Biased Technical Change and Earnings in Levels versus Growth

Task-based models of technical change begin from the observation that machines do not replace jobs so much as they reallocate tasks. Autor, Levy, and Murnane show that computerization substituted for routine cognitive and manual tasks while complementing non-routine analytic work, raising the return to education while reshaping task demand [11]. Autor emphasizes that automation also expands output and shifts demand for tasks in which labor retains a comparative advantage, mitigating mass unemployment [12]. Acemoglu and Restrepo formalize this tension as a race: automation displaces labor from existing tasks and can reduce the labor share [13, 14]. Goldin and Katz contextualize the schooling gradient in a longer race between education and technology: when the supply of educated labor fails to keep up with skill demand, the college premium rises even if the underlying technology is not uniformly skill-replacing [16].

Exposure indices inherit that dual structure. Because better-educated workers sort into information-intensive occupations, a high score naturally coincides with high absolute pay in a single cross section. However, as AI capabilities advance, the same occupations may face sharper substitution. Consequently, a high baseline exposure can simultaneously predict weaker subsequent income growth while the cross-sectional association stays positive.

Recent evidence on generative AI underscores this complexity. Eloundou et al. show that the new capability frontier is greater in high-wage occupations [3]. Brynjolfsson, Li, and Raymond find the opposite pattern inside one firm: a generative assistant raised productivity more for less-experienced customer-support agents than for the most experienced workers [17]. A tool that boosts junior productivity can coexist with a market ranking that is higher in high-wage occupations.

2.3 AI Pilot Zones as Local Industrial Policy

Place-based industrial policy directs infrastructure, subsidies, and regulatory flexibility toward designated regions. China’s Ministry of Science and Technology began designating National New Generation AI Innovation and Development Pilot Zones in 2019 across selected municipalities, including Deqing County. Quasi-experimental firm studies of the rollout report gains in labor-investment efficiency, total factor productivity, and related innovation outcomes inside pilot areas [7, 8].

Evaluating such policy rollouts from a household-worker perspective introduces distinct methodological challenges. Public CFPS releases emphasize province codes [15, 9]. Assigning each province the year of its first pilot aggregates treated urban cores with untreated hinterlands. Where county identifiers match municipal names, city-level designations can be directly coded. Furthermore, pilot selection is justified: economically successful and digitally mature provinces are typically designated earlier. Because those provinces may host higher concentrations of high-exposure occupations, unobserved regional developments can also confound policy estimates. To address this, we incorporate 2019 provincial totals of the Peking University Digital Financial Inclusion Index and AI patents as controls [18].

Staggered policy rollouts complicate two-way fixed-effects estimation when treatment timing varies [19, 20]. Given that our two-wave window has only one post-period for newly designated provinces, we avoid mixing treated units into the comparison group by dropping “always-treated provinces” and evaluating a clean comparison between the newly designated “switcher” provinces and “never-treated” provinces. Finally, to account for measurement error in aggregate codes, we implement a leave-out occupation Bartik, defined as the mean

2020 exposure in the respondent’s two-digit International Standard Classification of Occupations (ISCO) cell computed across all other provinces. Following Goldsmith-Pinkham, Sorkin, and Swift, we interpret this specification strictly as a within-cell robustness check [21].

2.4 Hypotheses

H1. Education is positively associated with individual AI exposure, and a college education $\times$ mid-career age (35–49) interaction raises exposure further. Mean exposure for that group significantly exceeds that of less-educated workers outside mid-career.

H2. Conditional on demographics and occupation-group controls, higher AI exposure is positively associated with log income in a cross section, yet higher 2020 AI exposure predicts smaller Δ log income from 2020 to 2022—reflecting task-based substitution within cognitive-intensive employment.

H3. In a clean comparison of newly designated provinces with never-treated provinces, the average designation effect on Δ log income is small relative to the occupational exposure gradient. Furthermore, the exposure penalty remains statistically significant after controlling for 2019 digital financial inclusion as well as AI patent stocks.

3 Data and Empirical Strategy

3.1 Data and Sample

We construct our primary estimation sample using the 2020 and 2022 waves of the China Family Panel Studies, linked at the person level [15, 9]. The survey window spans the early diffusion of large language models after GPT-3. Crucially, fieldwork for the 2022 wave was largely finished before the commercial launch and widespread adoption of ChatGPT and Deepseek in China. Estimation draws on six blocks of survey content: Demographics and the person link supply the matching backbone, schooling and occupation supply the two dimensions that H1 concerns, and residence and income supply the place controls and the outcome. Table 1 names each block and what it contributes.

Table 1: Survey content used in matching and estimation.

Block	Content used
Person link	Respondent identifier used to match the 2020 and 2022 waves
Demographics	Sex, age, and employment status
Education	Highest completed schooling level
Geography	Province of residence and urban or rural status in each wave
Income	Two reported annual income items
Occupation	Occupation group and the four-digit ISCO occupation code

Six filters build the pooled levels sample, applied in the order shown in Table 2. The age window and the requirement of a usable occupation code account for most of the loss. The final step keeps only respondents whose occupation code matches an exposure cell, which removes a further 9,155 observations.

Table 2: Sample filters for the pooled levels construction.

Step	Filter	Initial → Remaining	Dropped
Age	Age between 18 and 65	55,531 → 40,374	15,157
Occupation code	Four-digit ISCO code reported	40,374 → 31,902	8,472
Sex	Sex recorded	31,902 → 31,902	0
Marital status	Marital status recorded	31,902 → 30,628	1,274
Education	Schooling level within the survey scale	30,628 → 30,628	0
Employment	Employed or self-employed	30,628 → 30,239	389
Exposure match	Occupation code matched to an exposure cell	30,239 → 21,084	9,155

This procedure yields a pooled levels sample of $N = 21,084$ observations with a non-missing exposure score. Estimating cross-sectional income models then restricts the sample to respondents with positive reported earnings ($N = 16,021$). Growth designs merge the two waves on the person link and lock 2020 exposure as the baseline. Requiring positive labor income in both waves produces a balanced income first-difference panel of $N \approx 5,456$. The employment first-difference sample contains $N = 8,076$ observations. Finally, the difference-in-differences sample further drops always-treated provinces and leaves a clean analysis sample of $N = 3,420$.

Growth and policy designs need constructed outcomes and treatments rather than raw survey items. Table 3 defines each construct and states where it enters.

Table 3: Constructed outcomes, treatments, and controls.

Construct	Definition	Where it enters
$\Delta \log \text{ income}$	$\log(\text{income}_{2022}) - \log(\text{income}_{2020})$	First-difference and DID outcome
$\Delta \text{ employment}$	Employment status in 2022 minus 2020	First-difference outcome
College-educated, ages 35–49	Schooling at or above the college threshold and age 35–49 in 2020	H1 specification
Newly designated province	Province first treated after 2020 and by 2022	DID treatment
Leave-out occupation exposure	Mean exposure in the same two-digit ISCO cell, excluding the respondent's own province	Measurement-error check
2019 log digital finance index	Log of the provincial Digital Financial Inclusion Index total	Provincial control
2019 log AI patents	Log of one plus the provincial AI patent count	Provincial control

We merge two province panels locked at the 2020 province of residence. The first is the official 2019 provincial total of the Peking University Digital Financial Inclusion Index [18]. The second is city-level AI patent counts from 1985 to 2024, aggregated to provinces; we use the 2019 count as a predetermined measure. Where CFPS county identifiers match city names, we also attach the city designation dates and city AI patent counts for geographic checks.

3.2 Measuring AI Exposure

Each survey respondent inherits a year-specific occupational score from the SOC-by-year table in Zhang et al. [4]. ISCO codes map to SOC through the IBS crosswalk [25]. Missing occupational codes are imputed using the prefix-average scores from codes sharing the same leading digits. In longitudinal growth specifications, we lock each worker's occupation and exposure score at their 2020 baseline relative to Δy , preventing the regressor

from updating due to job switching. As a check, we also attach the 2018 score to the same 2020 occupation and estimate a 2018–2020 growth placebo, reported with the growth results in Section 4. The married indicator equals one for respondents living with a spouse or a cohabiting partner.

Table 4 describes the resulting estimation variables. Mean exposure is 0.450 with a standard deviation of 0.192. The score ranges from 0.136 to 0.849, so occupations differ widely within one sample. The average respondent is 39.9 years old, 56.0 per cent are men, and 76.8 per cent live with a spouse or partner.

Table 4: Summary statistics for estimation variables, pooled levels sample.

Variable	<i>N</i>	Mean	Std	Min	Max
Age	21,084	39.87	11.32	18	65
Male	21,084	0.560	0.496	0	1
Education (0–10 scale)	21,084	4.56	1.66	0	10
Married	21,084	0.768	0.422	0	1
AI exposure	21,084	0.450	0.192	0.136	0.849

3.3 Policy Treatment: AI Pilot Zones

We assign each province the approval year of its first national AI Pilot Zone city, with Zhejiang coded based on Deqing County. Approvals in September or later are coded as effective the following calendar year. We then divide provinces into three cohorts:

1. **Provinces first treated by 2020** are always-treated in our two-wave window.
2. **Provinces first treated in 2021–2022** are newly designated.
3. **Provinces with no pilot by 2022** are never treated.

Each individual’s policy status is locked at their baseline 2020 province of residence to prevent instances of endogenous migration. Where CFPS county identifiers can be matched directly to municipal names, parallel city-level checks are coded directly.

3.4 Empirical Strategy

To test H1 (Demographic Allocation) and H2 (Levels Premium), we estimate cross-sectional OLS regressions:

Levels.

$$E_i = \alpha + \beta_1 \text{Edu}_i + \beta_2 \text{Age}_i + \beta_3 \text{Urban}_i + \mathbf{X}'_i \gamma + \varepsilon_i, \quad (1)$$

$$E_i = \alpha + \beta_1 \text{Edu}_i + \beta_2 \text{Age}_i + \beta_3 \text{Urban}_i + \delta H_i + \text{interactions} + \mathbf{X}'_i \gamma + \varepsilon_i, \quad (2)$$

$$Y_i = \alpha + \theta E_i + \mathbf{Z}'_i \psi + u_i. \quad (3)$$

Equations (1)–(2) speak to H1. Equation (3) is the cross-sectional side of H2.

First differences. To test H2 (Growth Penalty) and H3 (Task Vulnerability vs Policy), first-differencing eliminates time-invariant individual heterogeneity[26]:

$$\Delta y_i = \alpha + \pi E_{i,2020} + \mathbf{W}'_{i,2020} \lambda + v_i. \quad (4)$$

Staggered difference-in-differences. Staggered approval can mix already-treated units into the comparison group when treatment timing varies [19, 20]. With only two waves we drop always-treated provinces, compare

newly designated with never-treated provinces, and estimate

$$\Delta y_i = \alpha + \rho \text{Switcher}_i + \pi E_{i,2020} + \phi(E_{i,2020} \times \text{Switcher}_i) + \mathbf{W}'_{i,2020} \lambda + v_i. \quad (5)$$

Under H3, if earnings risk is task-driven, the occupational exposure gradient (π) will remain large and negative. On the other hand, the policy treatment effect (ρ) will remain imprecise.

Leave-out Bartik. To address classical measurement error in matching aggregate occupation codes, we reestimate equation (4) using a leave-out Bartik instrument, defined as the mean exposure in the respondent's two-digit ISCO cell, leaving out the respondent's own province.

$$\Delta \log \text{income}_i = \alpha + \pi E_{i,2020} + \mathbf{W}'_{i,2020} \lambda + v_i, \quad (6)$$

instrumenting $E_{i,2020}$ with Z_i . Because two-digit cells dominate individual scores, the first stage is nearly mechanical. Following the share-identification reading of Bartik designs [21], we treat the exercise as a classical measurement-error check within occupation cells, not as an exclusion restriction against national occupation demand shocks.

Provincial controls. We re-estimate the income first-difference with the log of the 2019 provincial Digital Financial Inclusion Index and $\log(1 + \text{AI patents})_{p,2019}$ as additive controls, where p is the locked 2020 province.

Concurrent labor-market shocks. We also re-estimate equation (4) with provincial COVID case intensity and industry-level real activity shocks. COVID intensity is $\log(1 + \text{year-end confirmed cases})$ in the locked 2020 province. Industry shocks are national NBS value-added and urban non-private employment growth relative to 2019, matched to the respondent's 2020 industry.

Threats to identification. Table 5 lists the identified threats to the three designs, the direction of the implied bias, and the response we adopt.

Table 5: Threats to identification and design responses.

Threat	Direction	Response
Selection into titles	High-ability workers sort into high- E jobs	Levels as association; first differences remove person intercepts
Non-random pilots	Richer provinces get pilots earlier	Drop always-treated; newly designated vs never-treated
Province aggregation	City pilots diluted to province	Lock 2020 province; city-level DID where matched
Time-varying industry shocks	Confound E	COVID intensity and NBS industry shocks vs. 2019
Measurement error in E	Attenuates π	Leave-out Bartik as measurement-error check
Score-year choice	2020 score vs earlier score	2018 score on 2020 occupation; 2018–2020 placebo FD
Local confounders	Digital provinces host high- E jobs	Control 2019 digital finance index and AI patents
Short pre-period	Two-wave timing only	State as design limit

The 2020 occupation and the 2020 Zhang score are locked relative to Δy . Pilot designation is staggered but selected. Province clustering widens interaction standard errors. The provincial controls use 2019 measures

relative to 2020–2022 income changes.

4 Results

4.1 Who Holds High Exposure Scores

Education is the leading covariate in Table 6, which estimates equations (1)–(2). An increment on the education scale is associated with 0.048 higher exposure. College-educated workers aged 35–49 have an additional 0.062. Both estimates are affirmed by $p < 0.01$.

Table 6: OLS determinants of individual AI exposure (HC1 robust SEs).

Variable	(1) Baseline	(2) Interactions	(3) + Year FE
Education	0.0476*** (0.0008)	0.0370*** (0.0012)	0.0368*** (0.0012)
Age	-0.0005*** (0.0001)	-0.0009*** (0.0001)	-0.0008*** (0.0001)
Urban	0.0269*** (0.0025)	0.0270*** (0.0025)	0.0276*** (0.0025)
Male	-0.0571*** (0.0024)	-0.0537*** (0.0024)	-0.0539*** (0.0024)
Married	-0.0142*** (0.0031)	-0.0142*** (0.0032)	-0.0149*** (0.0032)
Occupation group	-0.0392*** (0.0021)	-0.0370*** (0.0021)	-0.0373*** (0.0021)
Log income	0.0034*** (0.0004)	0.0025*** (0.0004)	0.0026*** (0.0004)
Education $\times$ Age		-0.0008*** (0.0001)	-0.0008*** (0.0001)
Education $\times$ Urban		0.0132*** (0.0014)	0.0134*** (0.0014)
College $\times$ ages 35–49		0.0616*** (0.0048)	0.0621*** (0.0048)
Survey year 2022			-0.0092*** (0.0023)
Constant	0.3964*** (0.0095)	0.4428*** (0.0101)	0.4485*** (0.0101)
N	21084	21084	21084
R^2	0.2381	0.2535	0.2540

Note: HC1 robust SEs. Married equals one for respondents living with a spouse or a cohabiting partner. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

College-educated workers aged 35–49 also have a higher raw mean, 0.601 against 0.404 for everyone else, plotted in Figure 1.




Figure 1: Mean AI exposure by college education and mid-career status. Error bars are standard errors of the mean.

The education slope and the college-by-mid-career term stay in the same range when the sample is split by wave or by urban and rural residence. Table 7 puts the education coefficient between 0.035 and 0.052 and the interaction between 0.054 and 0.083, each with $p < 0.01$. Appendix A adds a short nonlinear check of the same education–age surface.

Table 7: Robustness of exposure determinants by survey year and urban/rural subsample.

Variable	(1) 2020	(2) 2022	(3) Urban	(4) Rural
Education	0.0368*** (0.0017)	0.0374*** (0.0017)	0.0515*** (0.0011)	0.0351*** (0.0013)
Age	-0.0005*** (0.0002)	-0.0012*** (0.0002)	-0.0006*** (0.0002)	-0.0013*** (0.0002)
Urban	0.0276*** (0.0034)	0.0265*** (0.0037)		
Education × Age	-0.0009*** (0.0001)	-0.0007*** (0.0001)	-0.0008*** (0.0001)	-0.0008*** (0.0001)
Education × Urban	0.0116*** (0.0020)	0.0154*** (0.0020)		
College × ages 35–49	0.0692*** (0.0068)	0.0539*** (0.0068)	0.0558*** (0.0054)	0.0826*** (0.0116)
Constant	0.4367*** (0.0141)	0.4473*** (0.0143)	0.4006*** (0.0131)	0.4573*** (0.0145)
N	10646	10438	13125	7959
R^2	0.2570	0.2544	0.2458	0.2092

Note: HC1 SEs. Demographic controls include an indicator for living with a spouse or a cohabiting partner. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4.2 Pay in Levels, Slower Growth

On the 16,021 respondents with positive income, higher exposure entails higher log earnings (Table 8). The uncontrolled slope is 1.12. Demographics bring it down to 0.66, and occupation-group indicators take it only to 0.64. All three estimates have $p < 0.01$.

Table 8: Log income on AI exposure (levels, positive-income sample).

Variable	(1) Bivariate	(2) + Demographics	(3) + Occupation-group indicators
AI exposure	1.1164*** (0.0377)	0.6603*** (0.0438)	0.6384*** (0.0438)
Education		0.1482*** (0.0055)	0.1447*** (0.0055)
Age		-0.0025*** (0.0008)	-0.0022*** (0.0008)
Urban		0.1309*** (0.0149)	0.1249*** (0.0149)
Male		0.5034*** (0.0148)	0.5057*** (0.0148)
Married		0.2338*** (0.0193)	0.2326*** (0.0193)
Constant	9.9866*** (0.0184)	9.0489*** (0.0441)	8.7255*** (0.0855)
Occupation-group indicators	No	No	Yes
N	16021	16021	16021
R ²	0.0497	0.1792	0.1818

Note: Sample restricted to respondents with positive reported income ($N = 16,021$). Column (3) enters occupation-group indicators rather than a linear code. HC1 robust SEs. Married equals one for respondents living with a spouse or a cohabiting partner. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Figure 2 plots the same association in decile means for income and employment.




Figure 2: Decile means with 95% confidence intervals for log income, left, on the positive-income sample, and employment, right, on the full levels sample, against AI exposure. The dashed line is the unconditional fit, slope 1.12, matching column 1 of Table 8.

Income rises across the deciles. Employment does not: Table 9 shows only a small positive association. Exposed workers still earn more in a cross section.

Table 9: Employment on AI exposure (levels).

Variable	(1) Bivariate	(2) + Demographics	(3) + Occupation group
AI exposure	0.0290*** (0.0105)	0.0299** (0.0121)	0.0250** (0.0122)
Education		0.0046*** (0.0015)	0.0051*** (0.0015)
Age		-0.0039*** (0.0003)	-0.0039*** (0.0003)
Urban		-0.0138*** (0.0043)	-0.0142*** (0.0043)
Male		0.1242*** (0.0044)	0.1239*** (0.0044)
Married		0.0328*** (0.0056)	0.0315*** (0.0056)
Occupation group			-0.0078*** (0.0024)
Constant	0.8832*** (0.0053)	0.9301*** (0.0140)	0.9598*** (0.0165)
N	21084	21084	21084
R ²	0.0003	0.0557	0.0562

Note: HC1 robust SEs. Married equals one for respondents living with a spouse or a cohabiting partner. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Income first-differences, however, offer an alternative explanation. In the $N = 5,456$ growth sample of equation (4), baseline exposure is -0.146 without controls and about -0.25 after demographics ($p < 0.01$; Table 10). Occupation-group indicators leave the coefficient near -0.23 .

Table 10: First-difference: change in log income (2020→2022) on baseline AI exposure (HC1).

Variable	(1)	(2)	(3)	(4)
Exposure	-0.1464** (0.0630)	-0.2498*** (0.0749)	-0.2360*** (0.0797)	-0.2325*** (0.0798)
College, ages 35–49			-0.0969 (0.1249)	-0.0945 (0.1249)
Exposure × mid-career			0.0090 (0.2067)	0.0048 (0.2066)
Education		-0.0018 (0.0096)	0.0048 (0.0105)	0.0049 (0.0105)
Age		-0.0101*** (0.0013)	-0.0095*** (0.0014)	-0.0095*** (0.0014)
Male		0.0334 (0.0275)	0.0329 (0.0276)	0.0329 (0.0276)
Urban		-0.0000 (0.0057)	0.0001 (0.0057)	0.0001 (0.0057)
Constant	0.2710*** (0.0308)	0.6937*** (0.0807)	0.6426*** (0.0864)	0.5148** (0.2210)
N	5456	5456	5456	5456
R ²	0.0009	0.0144	0.0151	0.0152

Note: Sample requires positive income in both waves. HC1 robust SEs in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

One existing concern is that high-scoring occupations perpetually grow more slowly. Table 11 refutes this concern. Column 1 repeats the 2020–2022 demographic first difference. Column 2 keeps the 2020 occupation

but attaches Zhang's 2018 score; the coefficient is still near -0.24 . Column 3 uses 2018 exposure to predict 2018–2020 growth and obtains a positive coefficient.

Table 11: Income growth and exposure timing: 2020 score, 2018 score, and 2018–2020 placebo.

Variable	(1) 2020–2022, E2020	(2) 2020–2022, E2018	(3) 2018–2020 placebo
Exposure	-0.2498*** (0.0749)	-0.2441*** (0.0755)	0.2722*** (0.0643)
N	5456	5470	5452
R^2	0.0144	0.0139	0.0137

Note: Cols (1)–(2) outcome is $\Delta \log$ income 2020–2022 with 2020 demographics. Col (1) uses Zhang year-2020 SOC exposure matched to the 2020 occupation; col (2) attaches the Zhang year-2018 cell to the same 2020 occupation. Col (3) is $\Delta \log$ income 2018–2020 on 2018 exposure (pre-window placebo). HC1 SEs. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Provincial COVID intensity and industry value-added and employment growth relative to 2019 move the exposure coefficient only from -0.250 to about -0.231 (Table 12).

Table 12: Income growth on baseline exposure with provincial COVID intensity and industry shocks.

Variable	(1) Baseline	(2) + COVID	(3) + VA	(4) COVID+VA	(5) + Emp.
Exposure	-0.2498*** (0.0749)	-0.2476*** (0.0753)	-0.2311*** (0.0769)	-0.2301*** (0.0774)	-0.2309*** (0.0773)
Log COVID cases (2020)		-0.0057 (0.0116)		-0.0028 (0.0119)	-0.0038 (0.0120)
Industry VA growth vs 2019			-0.1554 (0.2222)	-0.1568 (0.2217)	-0.1971 (0.2470)
Industry emp. growth vs 2019					0.0601 (0.4031)
N	5456	5456	5170	5170	5137
R^2	0.0144	0.0144	0.0131	0.0131	0.0125

Note: Outcome is $\Delta \log$ income, 2020–2022. Demographics (education, age, gender, urban) included in all columns. COVID is $\log(1 + \text{year-end confirmed cases})$ in the locked 2020 province. Industry shocks are national NBS value-added and urban non-private employment growth relative to 2019, matched to the respondent's 2020 industry. HC1 SEs. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Employment first differences are small and the intervals are wide (Table 13).

Table 13: First-difference: change in employment (2020→2022) on baseline AI exposure (HC1).

Variable	(1)	(2)	(3)	(4)
Exposure	-0.0159 (0.0191)	-0.0249 (0.0226)	-0.0337 (0.0240)	-0.0306 (0.0242)
College, ages 35–49			-0.0278 (0.0413)	-0.0261 (0.0413)
Exposure × mid-career			0.0805 (0.0634)	0.0770 (0.0634)
Education		-0.0030 (0.0029)	-0.0040 (0.0031)	-0.0043 (0.0031)
Age		-0.0025*** (0.0004)	-0.0026*** (0.0004)	-0.0025*** (0.0004)
Male		0.0149* (0.0079)	0.0148* (0.0079)	0.0147* (0.0079)
Urban		0.0011 (0.0013)	0.0011 (0.0013)	0.0011 (0.0013)
Constant	-0.0152 (0.0094)	0.0941*** (0.0257)	0.1040*** (0.0275)	0.0881*** (0.0326)
N	8076	8076	8076	8076
R ²	0.0001	0.0076	0.0079	0.0080

Note: HC1 robust SEs in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The sharper income penalty is among workers already above median baseline exposure (Table 14). Urban and rural columns drop respondents with missing urban codes, so the two sample sizes sum to less than 5,456.

Table 14: Heterogeneity in income first-differences by baseline group (HC1).

Variable	College, ages 35–49	Other	Above-median exposure	Below-median exposure	Urban	Rural
Exposure	-0.2117 (0.1971)	-0.2392*** (0.0803)	-0.4136*** (0.1368)	0.0177 (0.3156)	-0.1919** (0.0935)	-0.3205** (0.1454)
Constant	0.0481 (0.5616)	0.6573*** (0.0875)	0.8210*** (0.1309)	0.5036*** (0.1428)	0.3408*** (0.0550)	0.6655*** (0.1375)
N	531	4925	2819	2637	3291	1775
R ²	0.0137	0.0134	0.0285	0.0063	0.0142	0.0133

Note: Each column is a subsample FD regression with demographic controls. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4.3 Occupational Exposure and Place-Based Policy

When always-treated provinces are dropped, newly designated provinces differ from never-treated ones by -0.035 in $\Delta \log$ income after demographics (Table 15). Once baseline exposure is included, that designation coefficient stays small and the exposure coefficient is -0.379 , $p < 0.01$. With the interaction it is still -0.300 , $p < 0.05$.

Table 15: Staggered AI Pilot Zone DID: Δ log income (switcher vs never-treated).

Variable	(1)	(2)	(3)	(4)
Switcher	-0.0590* (0.0331)	-0.0352 (0.0330)	-0.0321 (0.0329)	0.0408 (0.0794)
Exposure			-0.3794*** (0.0977)	-0.3002** (0.1275)
Exposure $\times$ switcher				-0.1635 (0.1678)
Education		-0.0150 (0.0104)	0.0071 (0.0122)	0.0069 (0.0122)
Age		-0.0102*** (0.0017)	-0.0102*** (0.0017)	-0.0102*** (0.0017)
Male		0.0566 (0.0354)	0.0341 (0.0355)	0.0344 (0.0356)
Urban		0.0057 (0.0080)	0.0058 (0.0080)	0.0057 (0.0080)
Constant	0.2558*** (0.0234)	0.6758*** (0.1018)	0.7474*** (0.1002)	0.7142*** (0.1066)
N	3420	3420	3420	3420
R^2	0.0009	0.0131	0.0172	0.0174

Note: Sample drops always-treated provinces (pilot effective by 2020). Switchers are provinces first treated between the 2020 and 2022 waves. HC1 SEs. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 16 shows why the restricted sample is thinner. Always-treated provinces, including early pilots such as Beijing, Shanghai, Guangdong, and Zhejiang, account for 2,036 of the 5,456 income first-difference observations. Their mean baseline exposure is 0.463, against 0.445 among those retained. Inside the restricted sample, newly designated provinces ($N = 1,648$) average 0.454 and never-treated provinces ($N = 1,772$) average 0.436. The exposure coefficient in Table 15 remains large in this thinner sample.

Table 16: Restricted difference-in-differences sample: mean baseline AI exposure by treatment group.

Sample	N	Mean exposure	Mean Δ log income
Full income FD	5456	0.4516	0.2049
Always-treated provinces, dropped	2036	0.4633	0.1671
Sample after dropping always-treated	3420	0.4446	0.2274
Newly designated provinces	1648	0.4542	0.1968
Never-treated provinces	1772	0.4356	0.2558

Note: Income first-difference subsample with non-missing demographics. Always-treated provinces are those with a pilot effective by 2020, among them Beijing, Shanghai, Guangdong, Zhejiang, and Tianjin. Exposure is the 2020 occupational score.

Employment difference-in-differences estimates are near zero (Table 17).

Table 17: Staggered AI Pilot Zone DID: Δ employment (switcher vs never-treated).

Variable	(1)	(2)	(3)	(4)
Switcher	0.0024 (0.0095)	0.0113 (0.0096)	0.0116 (0.0096)	0.0033 (0.0241)
Exposure			-0.0374 (0.0296)	-0.0466 (0.0374)
Exposure $\times$ switcher				0.0187 (0.0496)
Constant	-0.0203*** (0.0064)	0.1044*** (0.0311)	0.1128*** (0.0317)	0.1165*** (0.0334)
N	5176	5176	5176	5176
R^2	0.0000	0.0090	0.0094	0.0094

Note: Sample drops always-treated provinces (pilot effective by 2020). Switchers are provinces first treated between the 2020 and 2022 waves. HC1 SEs. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The full-sample intensity specification still has a negative exposure coefficient. The exposure $\times \Delta$ PilotPost interaction is about -0.25 under HC1, $p < 0.10$, and less precise when standard errors are clustered by province (Table 18).

Table 18: Intensity DID: baseline AI exposure $\times$ pilot-status switch and income growth.

Variable	(1) Pilot switch	(2) + Exposure	(3) Exposure $\times$ switch	(4) + College, ages 35–49
Δ PilotPost	0.0043 (0.0278)	0.0043 (0.0278)	0.1168* (0.0684)	0.1115 (0.0681)
Exposure		-0.2498*** (0.0749)	-0.1770** (0.0843)	-0.1662** (0.0846)
Exposure $\times \Delta$ PilotPost			-0.2484* (0.1413)	-0.2369* (0.1404)
College, ages 35–49				-0.0882** (0.0393)
Constant	0.6424*** (0.0809)	0.6930*** (0.0806)	0.6587*** (0.0839)	0.6109*** (0.0893)
N	5456	5456	5456	5456
R^2	0.0125	0.0144	0.0150	0.0156

Note: Full income FD sample. Δ PilotPost equals one for provinces that switch into AI-pilot status between waves (locked at 2020 province). HC1 SEs.

As a check for classical measurement error, Table 19 instruments own exposure using the leave-out mean exposure in the same two-digit ISCO cell. The first-stage coefficient is about 0.99. In this sample, the OLS estimate is -0.250 and the 2SLS estimate is -0.229 . The instrument is too close to the endogenous regressor to credibly satisfy an exclusion restriction against national occupation-demand shocks, so we keep it only as a within-cell robustness check.

Table 19: Leave-out occupation Bartik as measurement-error check: income growth on baseline AI exposure.

Variable	(1) First stage	(2) OLS	(3) IV
Leave-out occ. exposure (Bartik)	0.9912*** (0.0043)		
Exposure		-0.2498*** (0.0749)	-0.2291*** (0.0776)
N	5456	5456	5456
First-stage F (1 instr.)	52270.37		

Note: Instrument is mean 2020 AI exposure in the same 2-digit ISCO cell, leaving out the respondent's province. Columns (2)–(3) include demographic controls. Robust SEs. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The 2019 provincial controls do little to the first-difference slope (Table 20). Starting from about -0.250 , the log digital finance index leaves exposure at -0.240 , log AI patents at -0.238 , and both together at -0.238 , each with $p < 0.01$.

Table 20: Income growth on occupational exposure with 2019 digital finance and AI patent controls.

Variable	(1) Baseline E	(2) + digital finance	(3) + AI patents	(4) Both controls
Exposure	-0.2498*** (0.0749)	-0.2396*** (0.0753)	-0.2380*** (0.0754)	-0.2381*** (0.0754)
2019 log digital finance index		-0.2379* (0.1311)		0.0129 (0.2192)
Log AI patents (2019)			-0.0219** (0.0100)	-0.0227 (0.0168)
Constant	0.6937*** (0.0807)	2.0611*** (0.7570)	0.8887*** (0.1201)	0.8217 (1.1450)
N	5456	5456	5456	5456
R^2	0.0144	0.0150	0.0153	0.0153

Note: Income FD sample merged to the 2019 Peking University Digital Financial Inclusion Index and province-summed AI patents. The provincial measures are predetermined relative to 2020–2022 changes. HC1 SEs. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A city-coded version that uses matched 2020 city names and the designation dates of the eighteen pilot cities looks like the province design (Table 21). After demographics the designation coefficient is -0.035 in the province sample and -0.036 in the city sample. Adding baseline exposure leaves both near -0.03 , while the exposure coefficients are -0.379 and -0.387 , each with $p < 0.01$.

Table 21: Province- versus city-coded pilot DID for income growth.

Variable	(1) Prov. + demog.	(2) Prov. + E	(3) City + demog.	(4) City + E
Switcher	-0.0352 (0.0330)	-0.0321 (0.0329)	-0.0361 (0.0544)	-0.0333 (0.0540)
Exposure		-0.3794*** (0.0977)		-0.3870*** (0.0949)
N	3420	3420	3491	3491
R^2	0.0131	0.0172	0.0159	0.0205

Note: Columns (1)–(2) drop always-treated provinces. Columns (3)–(4) restrict to respondents with a matched 2020 city name and drop always-treated pilot cities. Switcher is locked to the 2020 location's pilot calendar. Demographics included throughout. HC1 SEs. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

With city identifiers for every respondent, treatment would follow those eighteen cities' designation dates on the full sample, and the always-treated units would be early pilot cities rather than whole provinces. In Table 21 the designation coefficient stays near -0.03 . The -0.30 to -0.38 figures in the restricted sample are

the occupational exposure coefficients. Province labels put the pilot city and the rest of the province on the same treatment. The designation coefficient is then closer to zero than the city-level contrast on a complete city file.

Figure 3 shows the same education gap in the raw distribution. Median exposure rises from 0.31 in the lowest education band to about 0.66 for college graduates. Age alone goes the other way: the median is 0.49 at ages 18–29 and 0.34 at ages 50–65, so the mid-career premium in Table 6 is a college-conditional difference, not an age peak. Urban medians are 0.47 and rural medians 0.34. Female medians are 0.48 and male medians 0.34, which matches the negative male coefficient in Table 6. Income quartiles follow schooling only loosely.




Figure 3: Distribution of AI exposure by age band, schooling, gender, income quartile, and urban versus rural residence. Boxes span the interquartile range, the inner line is the median, and the dashed line is the sample mean 0.450.

4.4 Summary of Findings

Schooling and mid-career status raise exposure. Higher exposure still coincides with higher pay after occupation-group controls. The same 2020 scores predict slower $\Delta \log$ income over 2020–2022, including when the 2018 score is attached to the 2020 occupation and after COVID and industry shocks. The 2018–2020 placebo does not reproduce the negative link. Average designation effects are small. Exposure coefficients remain large under both province and city coding, and the first-difference penalty is still there after the 2019 digital finance and AI patent controls.

5 Discussion and Policy

5.1 Policy Implications

Higher-educated workers in mid-career have higher exposure scores and still earn more in a cross section. The same 2020 scores predict slower income growth through 2022. That growth penalty remains after COVID and industry-shock controls, after a 2018 score is attached to the 2020 occupation, and after the 2019 digital finance index and AI patent counts are included. Newly designated pilot provinces show only a small average growth effect.

Targeting therefore has to use schooling, mid-career status, and occupational exposure together with place

labels. Wage insurance and mid-career retraining fit this group because exposed workers still have income to lose and already grow at a slower pace in this window [3, 6]. Career counseling can thus be added where household earnings, not a single employer's tool, are the object. A firm associated assistant can raise junior productivity even as occupation-level scores remain higher in high-wage jobs [17]. Low measured exposure among less-educated and rural workers can coexist with weaker later access to AI-complementary tasks [11, 13, 6].

5.2 Limitations

The public files constrain the design in three ways.

1. **Two-wave timing.** The primary growth design uses CFPS 2020 and 2022. The interval is long enough to observe a first-difference penalty, but it does not follow the same workers through later model releases or additional pilot cohorts. Additionally, locking the 2020 occupation and the 2020 Zhang score with respect to Δy does not pin down the job-ad text behind the scores before 2020. The 2018-score check and the 2018–2020 placebo are answers we have within the sample. A later CFPS wave would allow the same person-level design to run over a longer post period without altering the hypotheses.
2. **Province labels versus city pilots.** Pilot zones are named at the city or county level. The main difference-in-differences design assigns each province the year of its first pilot and therefore mixes treated urban cores with untreated hinterlands. Early designation is also selected. Dropping provinces that are always treated, comparing newly designated with provinces never treated, and reporting a city-level comparison where county identifiers match are what the public files allow. Restricted county files, or a city-level household survey, would let treatment follow the official map on the full sample. The designation coefficient under province coding is about -0.03 , closer to zero than that city-level contrast. The -0.30 to -0.38 coefficients are occupational exposure in the restricted sample.
3. **Occupational scores and informal work.** Exposure inherits SOC-to-ISCO crosswalks and leading-digit averages when fine codes are missing, and CFPS underrepresents gig and informal jobs [4, 25, 9]. Provincial AI patent totals also dilute hubs within provinces. A China-specific task diary, or an occupation table built from domestic job ads rather than a mapped SOC cell, would tighten the worker-level score without rewriting H1–H3.

6 Conclusion

Higher-educated workers in the middle of their career have a higher occupational AI exposure in CFPS 2020–2022, and they still earn more in a cross section. Education increases exposure by about 0.048 for each additional category on the survey scale. College-educated workers aged 35–49 have a further premium near 0.062. Group means are about 0.60 versus 0.40. Among workers with positive income, the same exposure still associates with about 0.64 higher log income after demographics and occupation-group indicators.

Nevertheless, baseline exposure predicts about 0.23 to 0.25 lower Δ log income over 2020–2022. Attaching the 2018 score to the same 2020 occupation leaves a penalty near -0.24 . A 2018–2020 growth regression does not reproduce the negative link. The newly designated pilot provinces show only a small average growth effect. In the sample that excludes always-treated provinces, the exposure coefficient remains approximately -0.30 to -0.38 . After the 2019 digital finance and AI patent controls, the first-difference penalty is still about -0.24 . Changes in employment are small.

Therefore, job-title rankings and provincial pilot maps miss the workers whose 2020 scores already forecast slower growth, even though those workers still report income. The wage insurance, retraining, and counseling that follow only the place label will also miss them. Adding schooling, mid-career age, and occupational score is

a more direct screen. None of this implies that less-educated or rural workers are safe. Low measured exposure can coexist with weaker later access to complementary tasks.

A later CFPS wave would lengthen the post period after wider adoption and after later pilot cohorts. Finer geography would let treatments follow city designation dates for the full sample. A domestic task table would reduce the reliance on mapped SOC cells and on formal-job coverage. Until those files arrive, the 2020–2022 household evidence supports a narrower claim: in this early window, mid-career earnings growth tracks occupational AI exposure more closely than it tracks which province was named a pilot.

A Additional Results: Nonlinear Predictors of AI Exposure

This appendix checks whether nonlinear models recover the same education–age allocation of exposure as H1 in the levels sample. The identification of growth in the main text is based on the first differences, the pilot comparison, and the provincial controls for 2019. The holdout is an 80/20 person-level split of the CFPS levels sample with occupation and exposure-matching codes removed from X .

A.1 Models and Fit

We compare Elastic Net, Random Forest, and XGBoost on demographics only [22, 23, 24]. Figure 4 reports holdout fit on the same 4,217 held-out respondents. R^2 is 0.320 for XGBoost, 0.314 for Random Forest, and 0.264 for Elastic Net, with MSE of 0.025, 0.025, and 0.027. Demographics alone therefore explain roughly one-quarter to one-third of holdout exposure variance. That is enough to trace the schooling gradient. It is not enough to replace occupation codes.




Figure 4: Holdout R^2 and MSE for nonlinear predictors of AI exposure, demographics only.

Figure 5 shows the Elastic Net path: schooling remains the leading covariate.




Figure 5: Polynomial Elastic Net coefficient paths, left, and signed standardized effects, right.

A.2 Age and Education

Figure 6 evaluates predicted exposure by age and education on the holdout. Both tree models stay flat and low across ages when schooling is low, and both put the high values at high schooling. The schooling gap also widens with age: mean XGBoost predictions fall from 0.44 to 0.38 from ages 18–29 to 50–65 for workers below the college cut, while rising from 0.58 to 0.62 above it. That widening is the pattern the OLS interaction in H1 picks up. It is not an age peak in either surface.




Figure 6: XGBoost and Random Forest predicted exposure by age and education on the holdout sample.

These checks describe allocation in the levels sample. They do not replace the growth designs in Section 4.

References

- [1] Carl Benedikt Frey and Michael A. Osborne. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254–280, 2017.
- [2] Edward Felten, Manav Raj, and Robert Seamans. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217, 2021.
- [3] Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308, 2024.
- [4] Dandan Zhang, Hang Yu, Lixing Li, Jiayin Hu, Yiqing Mo, and Hongbo Li. The measurement of AI exposure and its impact on labor demand in China: Evidence from large language models. *Management World*, 41(7):59–75, 2025.
- [5] Daron Acemoglu, David Autor, Jonathon Hazell, and Pascual Restrepo. Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340, 2022.
- [6] David J. Deming. The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4):1593–1640, 2017.
- [7] Zhi Wang, Mei Xue, and Yunlei Luo. How does the national new generation AI innovation and development pilot zone impact corporate labor investment efficiency? evidence from China. *SAGE Open*, 15(4):21582440251408488, 2025.
- [8] Weihua Zhou and Chuanjie Zhu. Artificial intelligence innovation and enterprise total factor productivity: Evidence from China’s national new generation of artificial intelligence innovation and development pilot zones. *Journal of Industrial Technological Economics*, 44(10):79–88, 2025.
- [9] Institute of Social Science Survey, Peking University. China Family Panel Studies (CFPS) Public Data, 2022. Data waves 2020 and 2022.
- [10] Songül Tolan, Annarosa Pesole, Fernando Martínez-Plumed, Enrique Fernández-Macías, José Hernández-Orallo, and Emilia Gómez. Measuring the occupational impact of AI: Tasks, cognitive abilities and AI benchmarks. *Journal of Artificial Intelligence Research*, 71:191–236, 2021.
- [11] David H. Autor, Frank Levy, and Richard J. Murnane. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333, 2003.
- [12] David H. Autor. Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30, 2015.
- [13] Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018.
- [14] Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019.
- [15] Yu Xie and Jingwei Hu. An introduction to the China Family Panel Studies (CFPS). *Chinese Sociological Review*, 47(1):3–29, 2014.

- [16] Claudia Goldin and Lawrence F. Katz. *The Race between Education and Technology*. Harvard University Press, Cambridge, MA, 2008.
- [17] Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. Generative AI at work. *The Quarterly Journal of Economics*, 140(2):889–942, 2025.
- [18] Feng Guo, Jingyi Wang, Fang Wang, Tao Kong, Xun Zhang, and Zhiyun Cheng. Measuring China’s digital financial inclusion: Index construction and spatial characteristics. *China Economic Quarterly*, 19(4):1401–1418, 2020.
- [19] Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277, 2021.
- [20] Brantly Callaway and Pedro H. C. Sant’Anna. Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230, 2021.
- [21] Paul Goldsmith-Pinkham, Isaac Sorkin, and Henry Swift. Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624, 2020.
- [22] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.
- [23] Leo Breiman. Random forests. *Machine Learning*, 45:5–32, 2001.
- [24] Jerome H. Friedman. Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232, 2001.
- [25] Wojciech Hardy, Roma Keister, and Piotr Lewandowski. Educational upgrading, structural change and the task composition of jobs in Europe. *Economics of Transition*, 26(2):201–231, 2018.
- [26] Joshua D. Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press, Princeton, NJ, 2009.

Acknowledgment

This project studies mid-career earnings risk when occupational AI exposure scores are attached to Chinese household workers, and how that risk compares with staggered designation of AI Pilot Zones. The empirical core uses CFPS 2020–2022 person files, year-specific occupational exposure scores from Zhang et al., the Ministry of Science and Technology rollout dates, and predetermined 2019 provincial digital finance and AI patent measures.

The paper has two authors. Daniel proposed the worker-side earnings-risk cut, cleaned the CFPS 2020 and 2022 person files, built the ISCO–SOC–exposure crosswalk, locked 2020 occupations for the growth designs, proposed the three neural network solutions, and drafted the English report. Robert coded newly designated provinces and the married indicator, merged the 2019 provincial digital finance and AI patent controls, ran the levels OLS, income first-differences, restricted difference-in-differences, COVID and industry checks, and the 2018-score and 2018–2020 placebo in Python with statsmodels, and edited the draft. The instructor crucially guided the overall direction and engaged in neither the coding nor the drafting of the paper. In addition, the instructor contributed to the Acknowledgment section of this paper. We pay our greatest thanks to our instructor Tianzhou He.

The paramount methodological difficulty we encountered during the drafting of this paper is the geographical data granularity. Since public CFPS release province level identifications rather than municipal codes, we struggled to precisely map between the two. We ultimately yielded by leveraging ai crosswalk when codes are effectively absent. This problem leads into the SOC-ISOC crosswalk as well, leading digits imputations required a solution we struggled to find.

AI was used. Gemini 3.5 Flash averaged missing ISCO prefix cells in the exposure crosswalk when fine codes were absent. ChatGPT GPT-5.5 and Cursor Composer 2 were used to check English grammar on completed drafts. AI was used for prefix averaging in the crosswalk, grammar checks, literature-search clues. The authors chose the hypotheses, designed the difference-in-differences sample, ran the regressions, and wrote the text. Table numbers come from the estimation code. Sample definition, all empirical estimates, figure generation from code, identification claims, and the interpretation of the 2018–2020 placebo were completed by the authors.