Automation and Teenage School Dropout: Evidence from China

Ruilong Li*

Department of Economics, The Chinese University of Hong Kong

July 12, 2026

Abstract

Automation reallocates labor demand across education groups, but the groups themselves are not fixed across generations: teenagers decide whether to stay in school in response to the new earnings structure. The direct evidence to date sits mainly at the college margin and generally finds that automation induces more education. This paper moves the same question upstream, to the pre-college stages where progression is still rationed at a gateway. Using the four Chinese census and micro-census waves of 2000–2015 and a shift-share instrument built from European sector-level robot diffusion, I find that a one-standard-deviation increase in prefecture robot exposure lowers the in-school rate by 0.42 percentage points at primary ages (7–11), 1.07 points at junior-high ages (12–14), and 3.75 points at senior-high ages (15–17); the age-by-age profile steepens systematically as children approach the junior-to-senior-high gateway, and the 2000–2015 endpoint long difference reproduces the same ordering. This is not a reversal of the schooling expansion: the senior-high-age in-school rate rose from 55.1% to 81.8% over the same period, and the long difference implies that a prefecture with a one-standard-deviation larger increase in exposure forwent 6.53 points of that rise. Two cuts tie the gradient to progression access. Measured by provincial university admission chances on the eve of the 1999 higher-education expansion (1998), the school-leaving response concentrates in provinces where the route upward was narrow and is close to zero where it was wide; at the household level, the negative response at the 15–17 gateway comes almost entirely from teenagers in households with no member holding a bachelor's degree or above, and the rural-hukou point estimates are more negative as well.

Keywords: industrial robots; automation; human capital; in-school rate; pre-college schooling; progression access; Bartik instrument; China.

*I gratefully thank participants at the CUHK-Shenzhen Summer School for valuable discussions, and Julius Ilciukas, Yifan Zhang, and Junsen Zhang for their insightful comments and generous feedback. All errors are my own.

1 Introduction

How does automation reshape the human-capital supply of the next generation? The task-based literature on technological change has established that machines, software, and industrial robots, by taking over routine tasks, reallocate labor demand, employment, and wages across occupations and education groups (Autor et al., 2003; Acemoglu and Restrepo, 2018, 2020). Acemoglu and Restrepo (2022) in particular map task displacement into the wage structure across given demographic and education groups, and show that automation accounts for an important part of the rise in the US education wage gap. But education groups are not fixed across generations: teenagers can change whether they stay in school, and which educational track they enter, in response to the new earnings structure. Automation therefore changes not only the price of existing skills but potentially, endogenously, the future supply of skills. This paper studies that front-end margin—how robot exposure affects the in-school decisions of pre-college school-age children and teenagers, before they enter the labor market. In doing so it pushes the task-displacement framework from "the price of a given skill group" to "which skill group an individual enters," and takes the race between technology and education (Goldin and Katz, 2008) into an institutional setting where progression is rationed.

The direction of the schooling response is theoretically ambiguous. If robots depress low-skill wages and raise the skill premium, they lower the immediate opportunity cost of schooling and raise the relative return to more education, lengthening time in school. If the shock first erodes family income and liquidity, educational investment may instead fall. Moreover, in a staged, selective progression system, what matters is not only whether college wages are higher but whether the marginal student can reach college at all. Staying in secondary school is a sequential decision that preserves the option of continuing once more information about grades and progression is revealed (Stange, 2012); the value of that option depends on the tracks the student can actually reach. Where the route to college is wide, the force by which automation raises the relative value of abstract skills more easily translates into staying enrolled; where it is narrow, continuing in secondary school still mainly bets on the routine work available after high school, and that wage ladder is itself devalued by automation. The same robot shock can therefore push young people toward more education through opportunity costs and the skill premium, or push marginal students out through family constraints and a lower reachable continuation value. Which force dominates, and whether it changes with age and with the educational gateway, is the empirical question of this paper.

A recent set of studies estimates automation-induced educational adjustment directly. German youth shifted from vocational training toward colleges and universities as robot exposure rose (Dauth et al., 2021); robot exposure in the United States raised college enrollment and, in the longer run, degree attainment (Gunadi and Ryu, 2023). Di Giacomo and Lerch (2026) further find that the added enrollment concentrates in capacity-elastic community colleges, that graduation majors shift toward applied fields, and that the dominant channel is the fall in opportunity cost from shrinking low-skill job opportunities rather than a rising college wage premium. Branco et al. (2025) find that individuals who grew up in more robot-exposed areas are more likely to complete a bachelor's degree and gain in relative earnings, with the skill premium and opportunity costs both at work. The published direct evidence thus broadly supports "automation induces more education." The conclusion is nonetheless sensitive to when exposure occurs: Kang (2026) compares robot exposure at

different points of the life cycle in US data and finds that pre-adulthood exposure temporarily lowers college enrollment, with affected individuals catching up only later.

This evidence comes almost entirely from advanced economies where progression is not rationed, and the opportunity-cost logic itself rests on an institutional premise: for the marginal young person, more education must be *reachable*. US colleges admit by application, and community colleges have low entry barriers and elastic capacity—which is exactly where the added enrollment lands (Di Giacomo and Lerch, 2026); in Germany, vocational training and the tertiary track are likewise permeable (Dauth et al., 2021). In such a system, when robots depress low-skill wages and raise the relative return to abstract skills, the marginal young person who wants more schooling can actually walk into college, giving up nothing more than a low-skill job. In most developing economies, by contrast, access to further schooling is not allocated by application but rationed by selective examinations and quotas: whether one continues depends on exam results and admission capacity, and the door to more education has a gate on it. This difference is not background that can be omitted; it determines whether a fall in opportunity cost is converted into more time in school. The narrower the gate, the more the option of continuing is staked on the routine work that secondary schooling has historically fed—precisely the wage ladder automation devalues most. Extrapolating the positive response of advanced economies directly to an environment where progression is rationed therefore has no obvious basis.

The institutional difference is equally stark in the numbers. When the 1999 expansion was launched, China's tertiary gross enrollment ratio was 6.6%, against 71.4% in the United States in the same year; by 2015, after the largest higher-education expansion in the world, China had risen to 47.1%—still under six-tenths of the US level (88.9%).[1] Nor is the gap specific to China: in 2015 the tertiary gross enrollment ratio averaged 31.0% in low- and middle-income economies against 76.2% in high-income economies; India (27.3%), Vietnam (29.7%), and Indonesia (32.5%) sit in the same range as China, while the United States ranks in the top tier even among advanced economies (OECD average 72.0%). In terms of whether the marginal young person who wants more schooling can actually walk into college, the United States is closer to the exception than to the rule; the answer this paper provides is therefore not only about China but about the majority of developing economies where college access is likewise rationed.

China also has direct evidence closely related to this paper. The closest is Liu et al. (2025). That working paper uses a 2005–2015 long-difference design across 336 prefectures, instruments prefecture robot adoption with contemporaneous sector-level robot adoption in other Asian countries, and finds that robot exposure significantly raises the actual in-school rates of local youth aged 16–18 and 19–22, attributing the result to a falling opportunity cost from contracting low-skill employment and to a rising skill premium; the same paper uses in-school rates at ages 7–12 and 13–15 as a placebo and reports insignificant estimates.

This paper differs from that work in estimand, identification, and sample coverage, and reaches a different conclusion. On the estimand: the residence-based in-school rate at ages 19–22 reflects both local progression

[1] Tertiary gross enrollment ratios (indicator SE.TER.ENRR) are from the World Bank's *World Development Indicators* (World Bank, 2026), defined as tertiary enrollment divided by the population within five years of secondary-school completion age; the ratio can therefore exceed 100%.

and the cross-prefecture migration that admission brings, so this paper stops at age 17 and keeps only the ages at which prefecture assignment is not yet determined by admission outcomes. On identification: the design follows Acemoglu and Restrepo (2020)—sector robot penetration weighted by base-period sector employment shares, instrumented with robot diffusion in other economies—adapted to the Chinese setting as in Giuntella et al. (2025), so the instrument is built from European sector-level robot diffusion rather than from Asian countries that sit on the same supply chain and are China's main source of industrial robots; the design uses a prefecture–wave–age panel across the four census waves of 2000–2015 with province-by-wave fixed effects, identifying from exposure differences within a province–wave rather than from a single-window cross-prefecture long difference.[2] On coverage: this paper covers all 7–17 year-olds in the prefecture, rural and urban alike, and estimates age by age rather than by coarse groups, which is what makes visible the shape of the response as it steepens toward the junior-to-senior-high gateway—the compulsory-age estimates here are not a placebo but the low-age end of that profile.

Rural teenagers are the focus of that comparison, not a detail of coverage. The option-value framework predicts the strongest response among students who are unlikely to reach college even if they continue: their option of staying in secondary school is staked mainly on the routine work available after high school, which is the wage ladder automation devalues most. In China, rural-hukou teenagers have long been in that position—the urban–rural gap in basic education is long-standing (Hannum, 1999), and the schooling progression of poor rural students has long been a target of intervention research (Loyalka et al., 2013; Yi et al., 2015). Consistent with that prediction, Section 6 shows that the rural slope is more negative in point estimate than the non-rural slope, and that the rural disadvantage concentrates at the *zhongkao* margin (though the rural–non-rural gap itself is imprecise). Setting the observation window at age 16 and above, or measuring the compulsory stages only in coarse age groups, therefore skips both the population and the gateway where the mechanism predicts the strongest response. This paper accordingly does not read its negative result as a general effect of automation in China, but as an important boundary on the positive college-enrollment evidence: when the object shifts upstream, to in-school decisions still constrained by a progression gateway, the educational response can differ.

China offers a sharp institutional setting for testing this proposition. After nine years of compulsory schooling, junior-high graduates are sorted by the *zhongkao* into academic high school, vocational high school, or work; senior high school lies outside compulsory education and carries direct costs, and college entry is then allocated by the *gaokao* and province-level quotas. The junior-to-senior-high transition is therefore the first post-compulsory gateway on the route to college. Information alone, or early commitments of financial aid, need not be enough to carry poor rural students across that gateway (Loyalka et al., 2013; Yi et al., 2015); the 1999 higher-education expansion, by contrast, raised upstream academic-high-school completion by widening downstream opportunity, with stronger effects where senior-high capacity was more ample (Zhang, 2026). These facts imply that access to higher education is not a background condition outside the automation shock but a state variable determining the option value of staying in secondary school. In 2000,

[2]The skill-premium mechanism in Liu et al. (2025) is estimated on the 2006–2014 Urban Household Survey, which covers neither agricultural-hukou households nor migrants; inference is clustered at the prefecture level, with no shock-level inference or Rotemberg weight diagnostics for the shift-share design. Both are reported in Section 5 and the appendix of this paper.

the in-school rate at senior-high ages (15–17) in this paper's sample was only 55.1%, which is precisely why this upstream margin had substantial room during the sample period.

This paper takes the question of robots and educational choice to the pre-college stages. I use the 2000 population census, the 2005 1% population survey, the 2010 census long form, and the 2015 1% population survey, collapsing the microdata into prefecture–census-year–age cells. Robot exposure follows the shift-share construction of Giuntella et al. (2025): Chinese sector-level robot adoption interacted with the prefecture's 2000 sector employment shares, instrumented with European sector-level robot diffusion interacted with the same base-year shares. The baseline controls prefecture, province-by-census-year, and age fixed effects, so identification comes from exposure differences across prefectures within the same province–census-year. All analysis is restricted to ages 7–17: from age 18 onward, college admission moves a student's residence—typically together with the hukou registration—to the prefecture of the admitting institution, and because the census records people at their place of residence, prefecture assignment becomes endogenous to the schooling outcome itself. The estimand is whether a pre-college teenager is *currently still in school*, not the local enrollment rate after age 18, which is jointly determined by migration and admission.

The results show a clear negative stage gradient between robot exposure and the in-school rate. A one-standard-deviation increase in robot exposure lowers the in-school rate by -0.42^{**} percentage points at primary ages (7–11), -1.07^{***} at junior-high ages (12–14), and -3.75^{***} at senior-high ages (15–17); the age-by-age estimates move from small declines at younger ages to a systematically steeper decline as children approach the senior-high gateway. The 2000–2015 endpoint long difference reproduces the same ordering (-0.66^{**}, -1.73^{***}, -6.53^{***}). This does not mean automation reversed China's contemporaneous schooling expansion: the senior-high-age in-school rate rose from 55.1% in 2000 to 81.8% in 2015, and the long-difference result means that a prefecture with a one-standard-deviation larger increase in exposure forwent 6.53 points of that rise. The senior-high slope survives the standard battery for shift-share designs—leave-one-sector-out estimates, pre-period characteristic trends, migration restrictions, alternative weighting, Rotemberg weight diagnostics, and shock-level inference—though the most conservative shock-level standard error counsels restraint about precision.

Two complementary cuts tie the stage gradient to progression access. At the province level, I use provincial university admission chances on the eve of the 1999 higher-education expansion (1998) to characterize the width of the route upward: the school-leaving response concentrates in provinces with lower admission chances and is close to zero where chances were higher. At the household level, households with and without a co-resident college-educated member respond almost identically through the compulsory-age range but split cleanly at the 15–17 gateway: the negative response comes almost entirely from teenagers in households with no member holding a bachelor's degree or above. The first cut varies the route upward across provinces; the second characterizes a family's attachment to that route; both are consistent with reachable progression opportunities cushioning automation-induced school-leaving. The census, however, observes neither test scores, subjective progression probabilities, family income shocks, nor expected wages, so these results are evidence of channel consistency rather than direct identification of the option-value mechanism.

The paper contributes to three literatures. First, it adds to the literature on robots and endogenous human-capital accumulation. Existing work studies vocational training, college enrollment, major choice, and final attainment (Dauth et al., 2021; Gunadi and Ryu, 2023; Di Giacomo and Lerch, 2026; Branco et al., 2025); this paper pushes the observation window forward to the pre-college schooling sequence at ages 7–17 and shows that the response steepens as age approaches the actual school-leaving gateway and attenuates where ex ante progression chances were wider. That evidence supplies an institutional boundary for the uniformly positive published results at the college margin, and extends Acemoglu and Restrepo (2022)'s analysis of wages within given education groups to the endogenous formation of education-group membership. Relative to the closest Chinese study (Liu et al., 2025), the increment is not to show once more that robots change education, but to identify the shape of the response further upstream and to put progression access at the center of the explanation.

Second, the paper connects the literature on local labor-demand shocks and schooling choices. Educational investment responds quickly to perceived returns, institutional returns, and immediate job opportunities (Jensen, 2010; Abramitzky and Lavy, 2014); contracting low-skill labor-market opportunities can raise high-school completion (Greenland and Lopresti, 2016), while expanding low-skill employment typically raises the opportunity cost of staying (Atkin, 2016; Charles et al., 2018). Economic opportunity from infrastructure, in turn, can raise educational investment through family resources and expected returns (Adukia et al., 2020). Evidence from China's trade shocks likewise shows that the educational response depends on the skill content of the added demand and on export destinations, with no mechanically consistent sign (Li, 2018; Zhang and Zhou, 2023). The variation this paper adds separates the "immediate exit wage" from the "terminal value of the tracks reachable after continuing in secondary school": even if robots pass through only weakly to a dropout's immediate wage, they can still reprice the routine work that secondary schooling has historically fed. Progression chances determine how far the marginal student can bypass that devalued path.

Third, the paper adds to the literature on educational sorting, the senior-high bottleneck, and higher-education expansion in China. Existing work shows that the expansion transmits upstream along the schooling sequence and raises secondary completion (Zhang, 2026). This paper identifies a market force acting on the same upstream margin: automation did not stop the 2000–2015 senior-high expansion, but more-exposed prefectures expanded more slowly; more importantly, that negative slope concentrates in provinces where the pre-expansion route to college was narrower. The paper thereby joins "downstream opportunity shapes upstream choices" from the education-capacity literature with "labor demand changes the return to education" from the technology-shock literature, and shows that the two are not independent channels: the educational consequences of a labor-market shock depend on whether the education system can deliver the marginal student to the skill tracks that benefit from the new technology.

Section 2 states the two-margin framework and its testable predictions; Section 3 describes data and design; Section 4 presents the age-by-age and stage results; Section 5 summarizes the identification tests; Section 6 reports heterogeneity by hukou, sex, admission chances, and household education; Section 7 concludes.

2 A Two-Margin Framework for the Schooling Exit Decision

Core intuition. Robots do not lower all returns to education proportionally; they stratify how hard different skill tracks are hit. Automation substitutes for codifiable, repetitive routine tasks, and the first casualty is the routine manufacturing wage ladder V^R that senior-high graduates enter; the college track V^H, built on abstract tasks, is hit far more lightly and may even appreciate; and teenagers who drop out enter the labor market before the shock materializes, so their immediate wage w does not move. Staying in school is an option written on future tracks: with probability π_i it leads to V^H, with probability $1-\pi_i$ it falls back to V^R. When progression chances are limited, the option is written mostly on the most-devalued routine track, and the robot shock depresses the student's continuation surplus; the higher is π_i, the more the option is written on the lightly-hit college track, and the less the same shock erodes that surplus. The framework thus delivers the two properties the paper uses throughout: within the range where progression chances are limited, the exposure effect is negative, and its absolute value decays monotonically in the progression chance π. The census observes only the extensive margin of being in school: it cannot show whether high-π_i students who stay enrolled study harder as the college premium $V^H - V^R$ widens.

The individual continuation decision. China's school system sorts by test score at every progression gateway. The *zhongkao* sorts junior-high graduates into academic high school, vocational high school, or exit; the *gaokao* then sorts survivors into bachelor's programs or shorter vocational-college tracks. Consider a student i in grade g deciding whether to continue. Continuing requires direct schooling costs c_g plus an individual net cost ε_i, and forgoes immediate employment; its payoff is an option: with probability π_i the student eventually enters higher education with value V^H; with probability $1-\pi_i$ the student stops at high school and enters employment at a high-school graduate's level, with value V^R. The value of exiting is w: leaving school now for work. The student continues if and only if

$$\pi_i V^H + (1-\pi_i)\, V^R - c_g - \varepsilon_i \;\geq\; w. \tag{1}$$

If $V^H > V^R$, the condition can be written directly as a *continuation threshold*:

$$\pi_i \;\geq\; \bar{\pi}_i(x) \equiv \frac{w(x) + c_g + \varepsilon_i - V^R(x)}{V^H(x) - V^R(x)}. \tag{2}$$

Students below $\bar{\pi}_i(x)$ leave school; students above it continue. Those far below the threshold had already left before the robot shock arrived, and their binary in-school indicator cannot fall further; those far above it still continue, and any additional study effort does not register in a binary in-school indicator. The observable in-school response therefore comes from students who sat near $\bar{\pi}_i(x)$ and are pushed across the continuation margin by the shock.

The exposure gradient of the robot shock. Let local robot exposure be x, and write $\partial V^R/\partial x = -\rho$ and $\partial V^H/\partial x = -\eta$. The alternative to schooling is entering the labor market *now*: the dropout enters at the going wage today, while exposure measures robot penetration that materializes gradually over the coming

years; the shock lands on the future tracks that stayers will enter, not on the immediate wage the dropout has already locked in, so I set $\partial w/\partial x = 0$ directly. The *exposure-gradient assumption* has only two parts:

$$\rho > 0, \qquad \eta < \rho. \tag{3}$$

The first comes from task composition: the routine manufacturing ladder that senior-high graduates enter is exactly the task set robots substitute for, and it devalues with exposure. The second requires only that the college track is hit more lightly than the routine ladder; the sign of η is unrestricted—robots may devalue, leave unchanged, or even raise the absolute value of the abstract-task college track. Because $\rho > \eta$, the college premium $V^H - V^R$ rises with exposure, preserving the "relative values are re-ranked" intuition. $\partial w/\partial x = 0$ is a substantive restriction that the data could falsify: Appendix A reintroduces pass-through to the dropout wage as a perturbation parameter and characterizes the robustness region; if the shock passed through heavily to dropout wages, the exposure slope would turn positive and return to the opportunity-cost logic of the US evidence.

Define student i's continuation surplus Δ_i as the left side of (1) minus the right side; then

$$\frac{\partial \Delta_i}{\partial x} = -\big[\pi_i\,\eta + (1-\pi_i)\,\rho\big] \;\equiv\; -L(\pi_i), \tag{4}$$

where the loading $L(\pi)$ weights the shock to the two legs of the enrollment option by the progression probability. The exposure-gradient assumption immediately yields two properties.

First, *direction: with limited progression chances, the exposure effect is negative.* $L(0) = \rho > 0$ and L is linear and decreasing in π, so $L(\pi) > 0$ if and only if $\pi < \pi^* \equiv \rho/(\rho - \eta)$. If $\eta \geq 0$ (the college track is devalued too), then $\pi^* \geq 1$ and the exposure effect is negative for every π; only when $\eta < 0$—robots raise the absolute value of the college track—and $\pi > \pi^*$ does the effect turn positive. That combination requires progression to be nearly certain, which is exactly the margin where the US evidence sits; in an environment where progression is rationed by selective examinations, the empirically relevant range is $\pi < \pi^*$.

Second, *decay: the higher is π, the shallower the shock.* $L'(\pi) = \eta - \rho < 0$: the higher is π, the smaller the weight $1 - \pi$ on the most-devalued V^R and the larger the weight on the lightly-hit V^H, so the same shock erodes the continuation surplus less. The exposure effect therefore decays monotonically in π.

For the surplus effect to become observable school-leaving, it must still pass through the selection threshold of (2): a negative surplus shock shows up as school-leaving only for students who sat near $\bar{\pi}_i(x)$ and whose surplus is pushed from non-negative to negative. Those far above the threshold—mostly high-π_i students—keep a positive surplus and stay enrolled; even if the widened college premium makes them prepare harder for the exams, the census's binary in-school variable simply records them as enrolled and cannot register that intensified response.

From individual choice to the in-school share. For a given progression-chance group, let ε_i have cdf F and density $f > 0$. The share of grade-g students who continue is

$$S_g \;=\; F\big(\pi V^H + (1-\pi)\, V^R - w - c_g\big), \tag{5}$$

and the marginal effect of robot exposure on the in-school share is

$$\frac{\partial S_g}{\partial x} = -f(\cdot)\, L(\pi), \tag{6}$$

strictly negative on $\pi < \pi^*$: the individual-level direction property passes through unchanged to the group in-school share. The decay property passes through as well. When the density is approximately uniform near the decision margin ($f' \approx 0$),

$$\frac{\partial^2 S_g}{\partial x\, \partial \pi} = f(\cdot)\,(\rho - \eta) \;>\; 0, \qquad \text{equivalently} \qquad \frac{\partial}{\partial \pi}\left|\frac{\partial S_g}{\partial x}\right| = -f(\cdot)\,(\rho-\eta) \;<\; 0:$$

the higher is π, the smaller the absolute exposure slope and the shallower the robot-induced depression of the in-school share (the full cross-partial including the f' term is in Appendix A).

The other comparative static used directly is the level effect of progression chances:

$$\frac{\partial S_g}{\partial \pi} \;=\; f(\cdot)\,\big(V^H - V^R\big) \;>\; 0,$$

whenever the college premium is positive: where progression chances expand faster (π rises more), the share choosing to continue into senior high school rises more—consistent with the evidence that the expansion raised academic-high-school completion and narrowed the urban–rural secondary gap (Zhang, 2026). Combined with (6), the framework yields two empirical predictions. First, in a sample environment where progression chances are limited, the exposure slope is negative at every chance level, and the negative response concentrates where progression chances are lower and more students sit near the continuation margin. Second, as π rises the slope decays toward zero—for the slope to turn positive would require robots to raise the absolute value of the college track with progression nearly certain ($\eta < 0$ and $\pi > \pi^*$), or heavy pass-through of the shock to the dropout's immediate wage (Appendix A); both cases are closer to the United States than to China in the sample period. An everywhere-negative exposure slope that decays with chances is therefore both a prediction and a test of the exposure-gradient assumption. The framework organizes the channels into testable shape predictions; it does not treat any single channel as a mechanism this paper has established. Section 6.3 tests the cross-region prediction; Section 6.4 tests the same corollary at the household level—splitting by whether a co-resident household member is college-educated—where the contrast is sharpest exactly at the senior-high gateway.

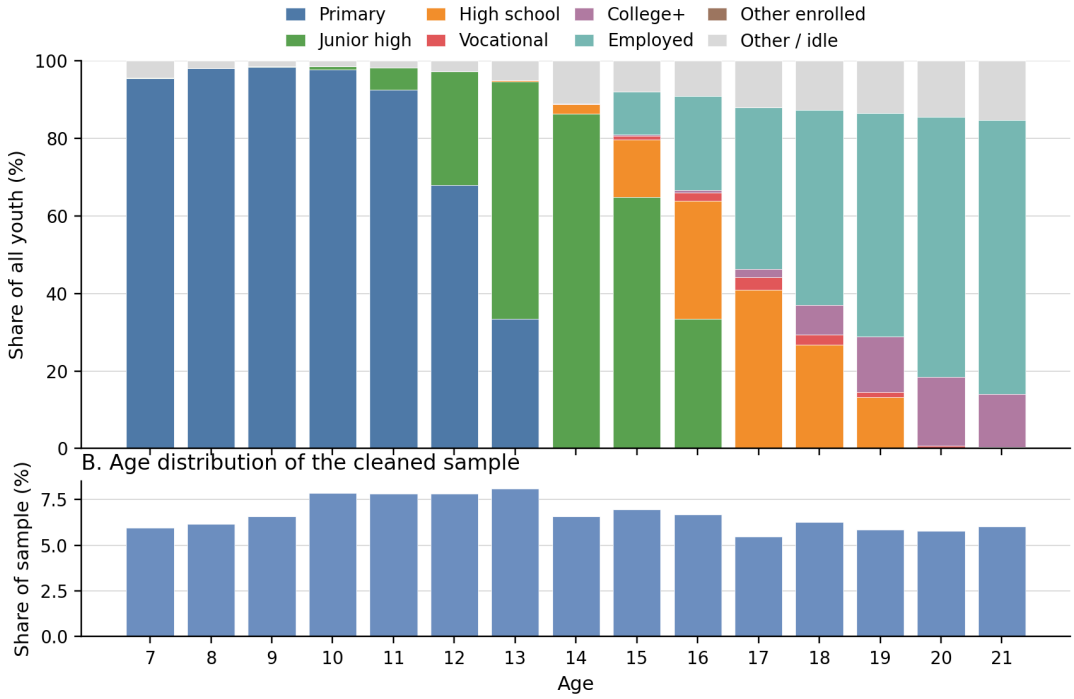
Two further extensions are developed in Appendix A: the extension to younger ages—an information option and the legal working age make the negative exposure slope steepen step by step as age approaches the senior-high gateway, the dropout ladder tested in Figure 4—and the cross-region implication. Section 6.3

tests the cross-region implication with provincial university admission chances on the eve of the expansion (1998).

3 Data and Empirical Strategy

Census samples. The main analysis uses the 2000 census micro-sample, the 2005 1% population survey, the 2010 long-form census, and the 2015 1% mini-census. The 1990 census provides pre-robot control variables. The sample drops cells that are clearly over-age for their grade (ages 14–15 in primary, 17 and above in junior high, 20 and above in academic high school). At the province level, the sample covers 26 of mainland China's 31 provincial units—all 22 provinces plus every prefecture-level unit of the four provincial-level municipalities of Beijing, Tianjin, Shanghai, and Chongqing—excluding only the five ethnic autonomous regions, whose education-policy environment (minority education policy, bilingual schooling, preferential admissions) is not comparable to the rest of the country. Each municipality contains only two prefecture-level units and contributes almost no identifying variation under province $\times$ wave fixed effects; including or excluding them leaves every stage estimate unchanged to the fifth decimal.

Age window: why the sample stops at 17. Figure 1 first shows the current school-level composition by age, which determines the paper's age ceiling. Ages 15–17 are the main carrier of the senior-high progression decision, with enrollees mostly in academic senior high or secondary vocational school, so the senior-high stage is described with ages 15–17. From age 18 onward, higher-education enrollees begin to appear among the cohort, and college admission moves a student's place of residence—typically together with the hukou registration—to the prefecture of the admitting institution. The census records people at their place of residence, so the prefecture assignment of people aged 18 and above is partly determined by the schooling outcome itself: prefectures with high progression rates "send away" their college-bound youth, while prefectures dense in colleges "import" students from elsewhere. Regressing the schooling or attainment outcomes of people aged 18+ on residence-prefecture exposure would therefore no longer measure local teenagers' schooling decisions, but would mix in this admission-induced compositional churn, with a bias of unknown sign whenever exposure correlates with the geography of colleges. To avoid this bias, the paper restricts all analysis to ages 17 and below—high school and below. Some senior-high students are still enrolled at age 18, so cutting the window at 17 is the conservative choice; outcomes at ages 18 and above would require linked data that can track individuals across prefectures, and are not examined in this paper.



Notes: Bars use all youth by age. Enrolled youth are classified by current school level; Employed is non-enrolled employed youth. The sample drops age-14/15 primary, age>=17 junior high, and age>=20 academic high-school cells.

Figure 1. Current school level, employment status, and age distribution. Descriptive figure, no estimating equation. The top panel plots $\Pr(\text{status} = s \mid age = a)$, where enrolled youth are classified by current school level and non-enrolled employed youth are shown as employed; the bottom panel plots $\Pr(age = a)$ in the cleaned analysis sample.

Cell-level specification. For the core results, I collapse each census outcome to prefecture–census-year–age cells, **pooled across hukou**, and relate the in-school level to the level of robot exposure within the four waves. For age a in prefecture p and wave $t \in \{2000, 2005, 2010, 2015\}$,

$$\bar{Y}_{p,t,a} = \beta \, \widetilde{\text{exp}}_{p,t} + \phi_p + \eta_{\pi(p),t} + \alpha_a + u_{p,t,a}, \tag{7}$$

where $\widetilde{\text{exp}}_{p,t}$ is China robot exposure standardized to unit variance, instrumented with the corresponding European shifter (both constructed in equation (8) below), and $\phi_p, \eta_{\pi(p),t}, \alpha_a$ are prefecture, province-by-wave, and age fixed effects. Identification comes from differences in robot exposure across prefectures within the same province-wave. I estimate (7) separately within each of three schooling stages—primary (ages 7–11), junior high (ages 12–14), and senior high (ages 15–17)—so β is that stage's levels in-school slope. The gradient of β across stages—steepening from primary through junior high to senior high—is the object of interest. The primary stage may also reflect budget, migration, and future-progression expectations; the paper's concern is not whether the primary slope equals zero, but whether the secondary stages steepen markedly relative to the primary stage. The main specification weights cells equally; the robustness section compares raw (population) cell counts and within-year normalized counts.

Bartik exposure. The exposure construction follows Acemoglu and Restrepo (2020), adapted to China as in Giuntella et al. (2025). For sector s, prefecture p, and census year t,

$$\begin{aligned} \exp_{p,t} &= \sum_s \ell_{p,s}^{2000} \frac{R_{s,t}}{L_{s,2000}} \times 1000, \\ \exp_{p,t}^{IV} &= \sum_s \ell_{p,s}^{2000} G_{s,t}^{EU}, \\ G_{s,t}^{EU} &= \frac{1}{|C_{s,t}|} \sum_{c \in C_{s,t}} \frac{R_{c,s,t}^{IFR}}{L_{c,s,t}^{EUKLEMS}} \times 1000, \end{aligned} \tag{8}$$

where $\ell_{p,s}^{2000}$ is prefecture p's 2000 sector employment share, $R_{s,t}$ is IFR's China robot stock, and $L_{s,2000}$ is 2000 sector employment. The IV shifter $G_{s,t}^{EU}$ is the average robot density by sector-year across the EU15 countries with positive EU KLEMS employment, excluding Luxembourg: IFR operational robot stock divided by EU KLEMS employment, in robots per thousand workers. I use China's current stock directly rather than subtracting the 2000 stock, because China's 2000 robot stock is approximately zero (930 robots installed nationwide); the paper therefore treats China's 2000 exposure as a zero-exposure pre-robot baseline, while the European shifter uses the observed 2000 European robot density.[3] IFR's early unspecified-sector manufacturing stock is allocated to sectors by identified sector shares, so 2005 is kept as a low-exposure wave. The levels coefficient is scaled to a one-standard-deviation change in the exposure **level** ($\widetilde{\exp}$), standardized to zero mean and unit variance on the four-wave estimation sample; the confirming long difference is scaled to a one-standard-deviation change in the five-year exposure **change**. Figure 2 shows the adoption timing and sector concentration: China is near zero in 2000, still low in 2005, accelerates after 2010, and concentrates in automobiles, electronics, plastics/chemicals, and metals.

Why the shifter is European. The exposure in equation (8) is built from China's own sector robot stocks, and Chinese adoption is itself endogenous: how many robots a prefecture installs depends on local labor costs, industrial policy, foreign investment, and product demand—the same forces that act directly on the schooling and employment opportunities teenagers face. Using local robot growth as the regressor would leave the estimated schooling response inseparable from those contemporaneous shocks. The instrument therefore retains only the advance of the technological frontier itself: robot penetration by sector in Europe reflects how far automation has become feasible in that sector worldwide—in which tasks robots have become cheap and reliable enough—which is determined by technical progress and European conditions, and is unrelated to the schooling decisions of teenagers in any Chinese prefecture. Multiplying the European sector shifter by the prefecture's 2000 sector employment shares therefore retains variation coming only from *how exposed the base-period industrial structure was to the global automation wave*, not from local demand for robots. Acemoglu and Restrepo (2020) instrument US exposure with a European shifter on exactly this reasoning, and Giuntella et al. (2025) carry the same logic to China. Asian countries that sit on the same supply chain and are China's main source of industrial robots are unsuitable as the shifter: the expansion of their robot

[3] Reconstructing exposure with the observed 2000 China sector stock (930 robots, allocated to sectors) rather than setting it to zero leaves the main results essentially unchanged: the senior-high levels estimate moves by less than 0.01 percentage point, and the 2000–2015 long difference is essentially unchanged.

industries is itself driven in part by Chinese demand, so exogeneity is hard to sustain. The cost is that the exclusion restriction falls entirely on the shares—the same 2000 sector shares must not affect schooling through any channel other than robots—which is what the remainder of this section and Section 5 address.

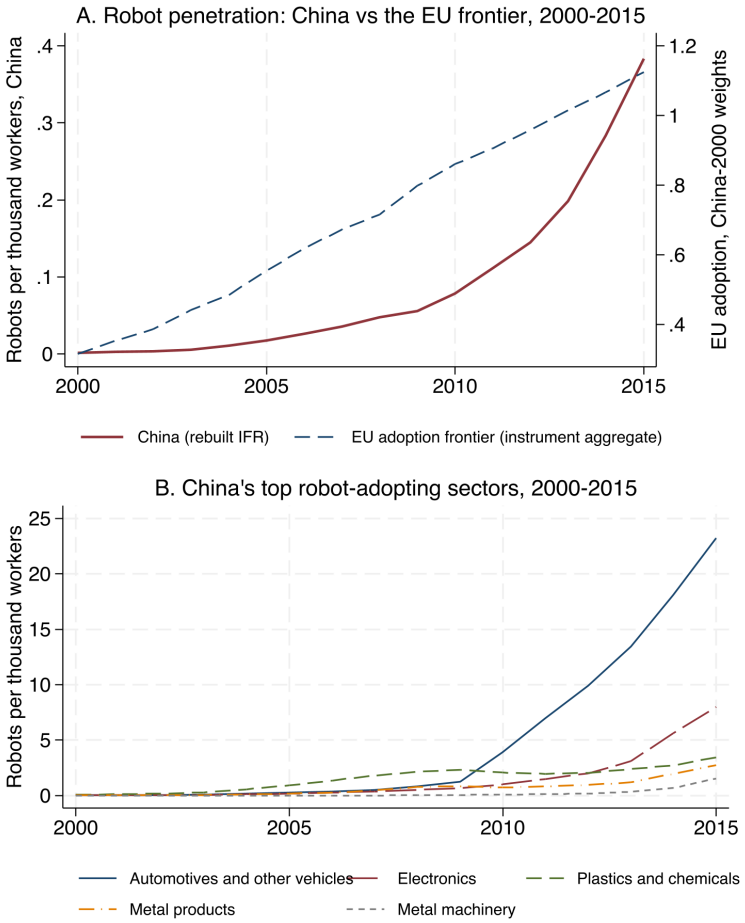


Figure 2. Industrial robot penetration, 2000–2015. Descriptive figure, no estimating equation. Panel A plots China robots per thousand 2000 workers and the EU adoption-frontier shifter aggregated with China's 2000 sector shares; Panel B plots sector robot penetration for the five top-adopting Chinese IFR sectors. No regression coefficient is plotted.

Long-difference estimand. To place the levels result alongside more conservative change-based evidence, I re-estimate the same stage gradient with a 2000–2015 endpoint long difference. Let $\Delta\bar{Y}_{p,a} \equiv \bar{Y}_{p,2015,a} - \bar{Y}_{p,2000,a}$, and let $\Delta\widetilde{\text{exp}}_p$ be the standardized change in robot exposure over the same interval. The long-difference regression relates $\Delta\bar{Y}_{p,a}$ to $\Delta\widetilde{\text{exp}}_p$, controlling for province and age fixed effects. This design tests the same pattern with fifteen-year changes rather than four-wave levels, so it is not another mechanism test but a conservative confirmation of the stage gradient.

The exclusion restriction is that the interaction of 2000 sector shares with European robot adoption affects schooling only through automation-related changes in the local opportunity set. The main levels specifi-

cation controls for prefecture, province-by-wave, and age fixed effects; the endpoint long-difference specification controls for province and age fixed effects, so both rely on differences in exposure across prefectures within a province. The remaining threat is that the same shares may proxy export demand, FDI, industrial upgrading, migration, or non-robot sector-specific shocks. Section 5 addresses these identification concerns directly. I construct the instrument from 2000 (rather than earlier 1990) sector shares because 2000 is the earliest wave with prefecture-by-IFR-sector employment that can build Bartik shares at the 18-sector level; the available 1990 census micro-data provide only aggregate prefecture sector shares (agriculture/manufacturing/secondary), insufficient to rebuild exposure at the 18-sector level. As an alternative, 1990 prefecture characteristics enter as $\times$ wave controls (Table 5, column (4)), absorbing differential trends correlated with 1990 industrial structure.

Pre-expansion admission chances. The provincial heterogeneity test (Section 6.3) uses a measure of progression-route width predetermined before the robot era: the 1998 provincial admission-chance rate, defined as provincial regular-college admissions divided by provincial regular senior-high graduates in the 1998 volume of the China Education Statistical Yearbook. 1998 is the last year before the 1999 expansion, and the expansion scaled admissions on each province's existing capacity, so the ratio captures both the pre-expansion width of the route and the room for post-expansion growth in college access. Provincial admissions sum to 1,083,627, digit-checked against the printed yearbook (the build script embeds anchor assertions). Two accounting caveats: admissions are counted at the institution's province, so large sending provinces (such as Shandong) read too low; and the distortion is most severe in the four municipalities—dense in universities, their ratios are severely inflated (up to 1.57 in Beijing) and do not measure local students' admission chances—so although they are inside the exposure analysis sample, they do not enter this section's chance grouping. Within the 22 provinces the ratio ranges from 0.19 (Qinghai) to 0.56 (Liaoning), and all splits (top/bottom quartile, above/below median) are computed within these 22 provinces. Figure 3 maps the spatial distribution of the ratio: the route upward is widest in the Northeast and in Shaanxi and Fujian, and narrowest across Shandong, Anhui, Guizhou, Gansu, and Qinghai; the juxtaposition of darker and lighter neighboring provinces provides the identifying contrast for Section 6.3.

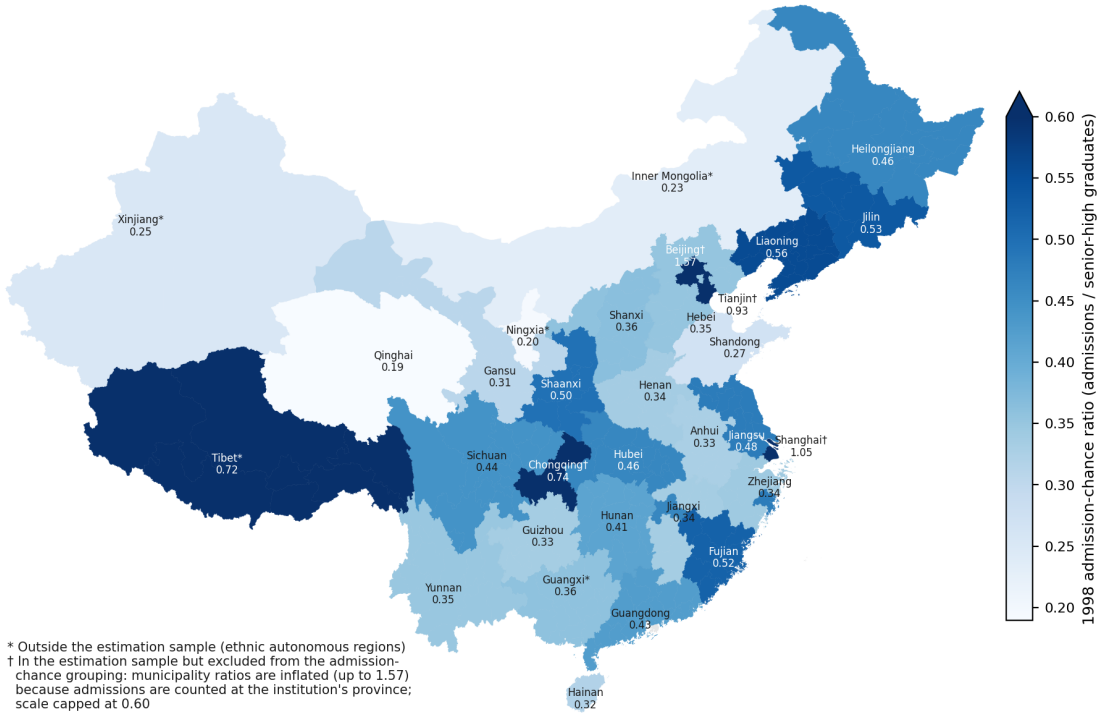


Figure 3. Provincial admission chances on the eve of the 1999 expansion. All 31 provincial units are shaded by the 1998 admission-chance ratio—regular-college admissions divided by regular senior-high graduates, both from the 1998 China Education Statistical Yearbook; darker means a wider route to university, and labels report each province and its ratio. Starred provinces (the five ethnic autonomous regions) are outside the estimation sample; daggered municipalities are inside the estimation sample but excluded from the admission-chance grouping, as their ratios are inflated—admissions are counted at the institution's province. The color scale is capped at 0.60.

Individual-level overview. Table 1 first summarizes the pooled sample at the individual level: roughly 3.11 million children aged 7–17 with a known hukou type across the four census waves. Nearly eighty percent hold rural hukou; the overall in-school rate is 88.6%, falling from 97.8% at primary ages to 66.0% at senior-high ages; among 15–17-year-olds, 24.6% are already employed. Among family-household children, only 2.5% live with a household member holding a bachelor's degree or above (6.8% at the junior-college-and-above definition)—the household heterogeneity in Section 6.4 splits on precisely this scarce trait.

Baseline schooling levels. Table 2 reports prefecture-year cell means pooled across hukou. Robot exposure is essentially zero in 2000, still low in 2005, and rises sharply by 2015. The table reports in-school rates for the three age stages, locating the relevant schooling margins before rural–non-rural heterogeneity is introduced below.

Table 1. Individual-level summary statistics, pooled 2000–2015 censuses. The sample is children aged 7–17 with nonmissing hukou type in the 2000, 2005, 2010, and 2015 census waves. Rows below the rule are computed on the family-household child file used in the household-education heterogeneity (collective households excluded; Section 6.4), so their sample size differs. Means and standard deviations are unweighted.

	Mean	SD	N
Age in years	11.98	3.00	3,113,188
Female	0.472	0.499	3,113,188
Rural hukou	0.783	0.412	3,113,188
Han ethnicity	0.928	0.259	3,113,188
Lives in a collective household	0.016	0.125	3,113,188
In school	0.886	0.317	3,113,188
In school, ages 7–11	0.978	0.147	1,408,094
In school, ages 12–14	0.939	0.239	922,725
In school, ages 15–17	0.660	0.474	782,369
Employed, ages 15–17	0.246	0.431	782,369
Household has a bachelor-or-above member	0.025	0.155	3,150,814
Household has a junior-college-or-above member	0.068	0.252	3,150,814
Household size	4.35	1.60	3,150,814

Table 2. Summary statistics by census wave. Each column summarizes prefecture–year cells pooled across hukou in that census wave. Schooling rows report in-school rates for ages 7–11, 12–14, 15–17, and 18–21; the analyses stop at age 17 because college admission migrates residence to the campus prefecture (Section 3).

	2000	2005	2010	2015
Robot exposure (raw)	0.000	0.014	0.064	0.309
Youth observations per prefecture-year	6,925.2	1,401.1	1,613.8	732.6
In-school rate, ages 7–11	0.970	0.973	0.983	0.955
In-school rate, ages 12–14	0.916	0.941	0.959	0.938
In-school rate, ages 15–17	0.551	0.715	0.785	0.818
In-school rate, ages 18–21	0.103	0.260	0.368	0.472
Prefecture-year cells	292	292	290	292
Underlying youth observations	2,022,156	409,116	468,009	213,907

4 Main Results: The In-School Decline Concentrates at the Senior-High Gateway

4.1 Levels evidence: the stage gradient

Age profile. Figure 4 first shows the pooled, age-by-age levels slope, each age estimated in its own regression. The profile is small across primary ages—about -0.6 percentage points at age 7—steepens age by age across secondary ages, and reaches a trough of about -5.4 points at age 17, the oldest age in the sample. This figure shows the overall shape first; Table 3 then aggregates ages into schooling stages. Group-specific responses appear in Section 6.

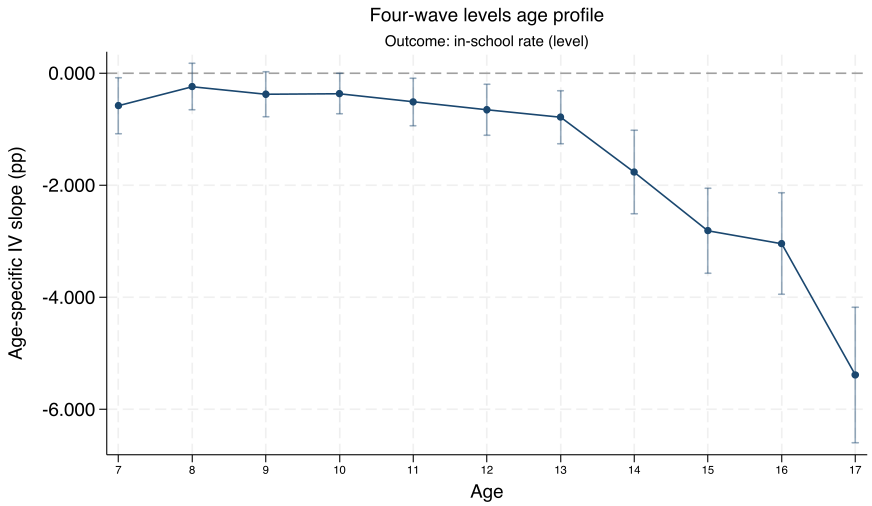


Figure 4. Robot exposure and in-school rate by age, pooled across hukou, in levels. Each age-a slope $\hat{\beta}_a$ is estimated in its own regression of the in-school rate on the exposure level, $\bar{Y}_{p,t,a} = \beta_a \exp_{p,t} + \phi_p + \eta_{\pi(p),t} + \varepsilon_{p,t,a}$, with prefecture and province-by-wave fixed effects, instrumented with the corresponding EU shifter. The plotted coefficient is $\hat{\beta}_a$, the age-a levels IV slope; whiskers are 95% CIs.

Stage-by-stage levels. Table 3 reports the four-wave levels version: on prefecture–census-year–age cells pooled across hukou, relating the in-school rate to current robot exposure and absorbing province-by-census-year fixed effects. Negative coefficients mean lower in-school rates in more-exposed prefectures. The stage gradient is clear: 2SLS is -0.42^{**} percentage points for primary (ages 7–11), -1.07^{***} for junior high (ages 12–14), widening to -3.75^{***} for senior high (ages 15–17). OLS gives the same stage ordering. The core evidence is precisely this primary-to-senior-high gradient.

This age gradient itself has economic content. Younger children are still far from immediate school-leaving and labor-market entry; compulsory-schooling constraints and academic information that has not yet been revealed keep most of them away from the continuation-versus-exit margin. As age rises, students learn more about their own performance and progression prospects, and leaving school for work becomes progressively more feasible. The same labor-market shock to the future returns to education and work is therefore more likely to translate into observable in-school changes at older ages. The *zhongkao*, the end of compulsory schooling, and the minimum working age provide the institutional backdrop for this process, but the paper does not interpret the data as identifying a separate institutional discontinuity at any single age. The more robust statement is: the closer teenagers are to the actual school-leaving margin, the stronger the negative relationship between robot exposure and the in-school rate.

Table 3. Stage-by-stage levels in-school effect (pooled), four waves. Each row is one schooling stage, estimated as a separate regression of the in-school rate on the robot-exposure level on prefecture × wave × age cells collapsed across hukou and pooled across the four census waves (2000, 2005, 2010, 2015). All specifications include prefecture, province-by-wave, and age fixed effects. Negative coefficients mean lower enrollment in more-exposed prefectures. 2SLS instruments the China exposure level with the European shifter level. Coefficients are percentage points per 1 SD exposure; SE are clustered by prefecture.

	Robot exposure (1 SD), pp			
Schooling stage	OLS	2SLS	First-stage F	Obs
Primary 7–11	-0.26*	-0.42**	527	5,828
	(0.14)	(0.17)		
Junior high 12–14	-0.59***	-1.07***	521	3,495
	(0.20)	(0.24)		
Senior high 15–17	-1.84***	-3.75***	521	3,497
	(0.34)	(0.39)		
Fixed effects	Prefecture, province × wave, and age FE			
Cell unit	Prefecture × wave × age (pooled across hukou)			
Waves	2000, 2005, 2010, and 2015			
Weights	Equal cells			

4.2 The 2000–2015 long-difference confirmation

Table 4 rewrites the same stage gradient as the 2015-relative-to-2000 change, controlling for province and age fixed effects. The differenced estimates reproduce the same ordering: -0.66^{**} percentage points for primary (ages 7–11), -1.73^{***} for junior high, and -6.53^{***} for senior high. This version no longer compares four-wave levels but directly asks, among prefectures where robot exposure rose more between 2000 and 2015, which age stages saw larger in-school declines.

Table 4. Cumulative 2000–2015 long-difference in-school effect (pooled). Each row is one schooling stage, estimated as a separate regression of the 2015-minus-2000 change in in-school status on the 2015-minus-2000 change in robot exposure, using prefecture × age cells collapsed across hukou. All specifications include province and age fixed effects. Negative coefficients mean enrollment fell more in prefectures with larger exposure changes. 2SLS instruments the China exposure change with the European shifter change. Coefficients are percentage points per 1 SD exposure change; SE are clustered by prefecture.

	Δ robot exposure, 2000–2015 (1 SD), pp			
Schooling stage	OLS	2SLS	First-stage F	Obs
Primary 7–11	-0.61**	-0.66**	463.7	1,320
	(0.24)	(0.26)		
Junior high 12–14	-1.46***	-1.73***	458.7	792
	(0.41)	(0.41)		
Senior high 15–17	-6.18***	-6.53***	458.7	792
	(0.66)	(0.70)		
Fixed effects	Province FE and age FE, within stage			
Cell unit	Prefecture × age (pooled across hukou)			
Difference	2015 minus 2000			
Weights	Equal cells			

Reading the two estimators together. The levels estimate is the paper's main description, capturing overall in-school differences in more-exposed areas across the four-wave sample; the 2000–2015 long-difference estimate uses only the endpoint change, testing whether the same stage gradient also appears in fifteen-year changes. What both support is not whether the primary stage is zero, but that the in-school slope steepens systematically as age approaches the senior-high gateway.

5 Identification Threats and Robustness

The exclusion restriction is that 2000 sector shares interacted with European robot adoption affect schooling only through automation-related changes in the local opportunity set. This section addresses the threats to that restriction in three groups: pre-existing differential trends correlated with 2000 sector geography (addressed with sector exclusions and predetermined-characteristic trends), composition and migration, and the mechanics of the shift-share instrument itself. The headline test appears in Table 5; the remaining tables are collected in Appendix B. Across all of them, the senior-high decline keeps its sign and approximate magnitude.

Sector exclusions and industrial trends. Table 5 directly examines the hukou-pooled senior-high-age in-school rate (ages 15–17). All columns absorb province-by-year fixed effects. Column (1) reproduces the senior-high slope of Table 3, -3.75^{***}. Columns (2)–(3) rebuild China exposure and the European instrument excluding electronics and automobiles respectively, testing whether the gradient is driven by a single flagship sector. Column (4) adds 1990 prefecture characteristics interacted with year. All seven characteristics come from the 1990 census micro-data and are determined before the arrival of robots: log population, the male share, the working-age (15–64) share, the college and secondary-and-above attainment stocks among those aged 25 and over, and the manufacturing and agricultural shares of employment; each is standardized across prefectures and interacted with year dummies, absorbing differential trends correlated with pre-period demographics, education stocks, and industrial structure. Column (5) adds 2000 high-school capacity interacted with year. Capacity is measured in the 2000 census as the share of a prefecture's residents aged 17–21 who have completed upper-secondary schooling (academic-high-school or *zhongzhuan* graduates, including those already in junior college or above), computed without splitting by hukou and standardized across prefectures, absorbing baseline differences in senior-high supply. The coefficient is negative under all specifications.

Table 5. Robustness of the pooled senior-high-age in-school slope, ages 15–17. Cells are prefecture × wave × age, pooled across hukou. Column (1) includes prefecture, province-by-year, and age FE; columns (2)–(3) rebuild exposure excluding electronics and excluding automotives; columns (4)–(5) add 1990 prefecture characteristics × year and 2000 high-school capacity × year. Negative coefficients mean lower in-school rates in more-exposed prefectures. 2SLS; prefecture-clustered SE; equal-cell weights; pp per 1 SD exposure level.

	(1)	(2)	(3)	(4)	(5)
In school, ages 15–17	-3.75***	-3.82***	-3.06***	-2.42***	-1.73***
	(0.39)	(0.53)	(0.46)	(0.48)	(0.47)
Observations	3,497	3,269	3,269	3,107	3,353
First-stage F	520.7	1747.1	173.8	907.4	390.0
Mean outcome	0.705	0.702	0.702	0.701	0.704
SD outcome	0.214	0.216	0.216	0.216	0.214
Prefecture and age FE	Yes	Yes	Yes	Yes	Yes
Province × year FE	Yes	Yes	Yes	Yes	Yes
Excluding electronics from exposure	No	Yes	No	No	No
Excluding automotives from exposure	No	No	Yes	No	No
1990 characteristics × year	No	No	No	Yes	No
2000 HS capacity × year	No	No	No	No	Yes
Weights			Equal cells		

Migration. Selective migration is the natural worry in a country of mobile households. Appendix Table A2 restricts the sample to youth whose five-year-prior residence was local or comparable within province: the primary-age response goes to near zero, while the in-school slope in the 15–17 senior-high window remains -3.62^{***} percentage points.

Instrument mechanics. Three diagnostics probe the shift-share construction itself. Rotemberg weights (Appendix Table A3) show the identifying variation concentrates where it should—automobiles plus electronics carry about 57.4% and the top five sectors 88.7%—so the exogeneity discussion above rightly centers on those sectors. Leave-one-sector-out estimates (Appendix Table A4) rebuild both exposure and instrument dropping each of the 18 IFR sectors in turn; the senior-high slope stays negative throughout. Cell-weighting choices do not change the conclusion either: the senior-high stage shows the largest decline under equal cells (-3.75), raw population counts (-7.05), and within-year normalized counts (-3.01) alike (Appendix Table A1).

Shock-level inference. Finally, Appendix Table A5 computes inference at the sector-shock level following Borusyak et al. (2022), addressing the concern raised by Adão et al. (2019). The Rotemberg decomposition and the shock-level inference now use exactly the same 15–17 five-year-change sample, standardization, and fixed effects, and both point estimates equal -2.91 percentage points. Sector-clustered shock-level standard errors still support statistical significance, but the most conservative heteroskedastic shock-level standard error is 1.93 points, leaving the estimate imprecise. The -6.53^{***} of Table 4 uses the 2000–2015 endpoint change and fifteen-year standardization, so the magnitudes are not directly comparable.

6 Who Carries the Gradient: Hukou and Sex

The previous results establish the reduced-form stage gradient. This section asks which groups carry it. The overall pattern is that the enrollment decline is shared across groups rather than concentrated in any one.

6.1 Hukou: a larger rural point estimate but an imprecise gap

To estimate the hukou difference, I reopen the levels cell design at the prefecture–census-year–hukou–age level, interacting both the exposure slope and the fixed effects with hukou, and estimate separately within each schooling stage:

$$\bar{Y}_{p,t,h,a} = \beta\, \widetilde{\exp}_{p,t} + \rho\, \left(\widetilde{\exp}_{p,t} \times \text{rural}_h\right) + \phi_{p\times h} + \eta_{\pi(p),t\times h} + \alpha_{a\times h} + u_{p,t,h,a}. \tag{9}$$

For each stage, β is the non-rural in-school slope and ρ is the rural-minus-non-rural gap. Figure 5 plots age-by-age levels slopes separately for rural (agricultural hukou) and non-rural (non-agricultural hukou) youth. Both profiles are negative across primary ages—the cross-sectional component falls on both groups—then separate at the senior-high entrance, with the rural slope below the non-rural one. The rural disadvantage concentrates at the junior-to-senior-high gateway—the *zhongkao* margin—rather than spreading uniformly across secondary ages.

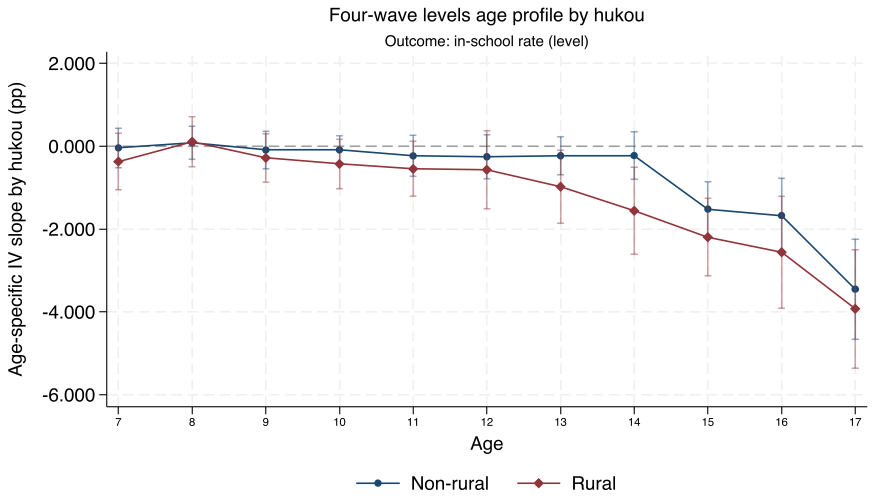


Figure 5. Robot exposure and in-school rate by age and hukou, in levels. Each hukou–age slope is estimated in its own regression of the in-school rate on the exposure level, $\bar{Y}_{p,t,a} = \beta_{h,a}\exp_{p,t} + \phi_p + \eta_{\pi(p),t} + \varepsilon_{p,t,a}$ within each hukou group, with prefecture and province-by-wave fixed effects, instrumented with the EU shifter; N denotes non-rural and R rural. Whiskers are 95% CIs.

Both the non-rural (-2.22^{***} percentage points) and rural (-2.88^{***} points) senior-high slopes are large and significant; the rural–non-rural gap is -0.66 points and insignificant. I therefore read the hukou profile as suggestive heterogeneity: the rural point estimate is more negative, but the most robust fact remains the overall stage gradient.

6.2 Sex: the enrollment decline is broadly symmetric

The schooling response is broadly symmetric across sex. Table 6 re-estimates the stage-by-stage levels design on prefecture–census-year–sex–age cells, interacting both the exposure slope and the fixed effects with sex. At senior high (ages 15–17), the male slope is -3.55^{***} percentage points and the female slope -4.01^{***}, with a female-minus-male gap of -0.46 points, insignificant. Figure 6 shows the same male–female overlap age by age.

Table 6. Stage-by-stage levels in-school effect by sex. Each column is a separate 2SLS regression for one schooling stage, using prefecture × wave × sex × age cells across the four census waves. Rows report the male slope, the female slope, and the female minus male gap. All specifications include prefecture × sex, province × wave × sex, and age × sex FE. Coefficients are percentage points per 1 SD exposure level; SE clustered by prefecture. Windows are 7–11, 12–14, and 15–17.

	2SLS: robot exposure (1 SD), pp		
	Primary 7–11 (1)	Junior high 12–14 (2)	Senior high 15–17 (3)
Male	-0.35*	-0.86***	-3.55***
	(0.18)	(0.23)	(0.41)
Female	-0.48***	-1.25***	-4.01***
	(0.18)	(0.30)	(0.46)
Female–male	-0.12	-0.39*	-0.46
	(0.14)	(0.21)	(0.36)
Observations	11,649	6,987	6,991
First-stage F	263	260	260
Fixed effects	Prefecture × sex, province × wave × sex, and age × sex FE		
Cell unit	Prefecture × wave × sex × age		
Waves	2000, 2005, 2010, and 2015		
Weights	Equal cells		

This symmetry is itself informative. The intuition that "automation should hit boys harder" builds on the male-dominated auto plants of Western labor markets; but the sectors carrying China's robot adoption in this period—besides automobiles, also electronics and textiles (Appendix Table A3)—place many young women on routine assembly lines. The "secondary-schooling–routine-work" pathway depressed by automation is therefore not narrowly male, and the enrollment response is correspondingly broad across sex rather than a narrow male-factory-work response.

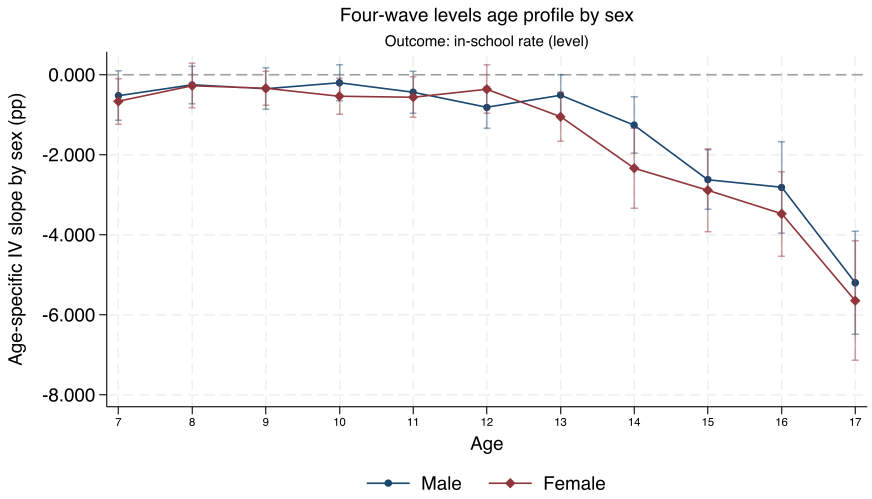


Figure 6. Robot exposure and in-school rate by age and sex, in levels. Each sex–age slope is estimated in its own regression of the in-school rate on the exposure level, $\bar{Y}_{p,t,a} = \beta_{g,a} \exp_{p,t} + \phi_p + \eta_{\pi(p),t} + \varepsilon_{p,t,a}$ within each sex, with prefecture and province-by-wave fixed effects, instrumented with the EU shifter; M denotes male and F female. Whiskers are 95% CIs.

6.3 Pre-expansion admission chances: weaker exposure responses where the route upward was wider

The higher-education expansion launched in 1999 scaled admissions on each province's existing capacity, so provinces with higher 1998 admission-chance rates also saw larger post-2000 growth in college access. In the language of Section 2, teenagers in those provinces faced a wider progression route (π higher, and rising further after the expansion): when the robot shock devalues the "secondary schooling–routine work" pathway (V^R falls), the left side of inequality (1) is eroded less in high-chance provinces, and more of the option value of staying enrolled survives. The prediction is that exposure-driven school-leaving should concentrate in low-chance provinces; in high-chance provinces the response decays toward zero rather than turning positive (the decay property of Section 2).

Table 7 splits the 22 provinces in the sample by the 1998 admission-chance rate of Section 3 into top and bottom quartiles (six provinces each; the middle half does not enter the contrast, and the four municipalities, whose chance measure is not comparable, do not enter the grouping), interacting the exposure slope with the group on prefecture–wave–age cells, with prefecture, province $\times$ wave, and age $\times$ group fixed effects. Panel A (no controls) shows deeper declines in low-chance provinces at every stage: pooled over ages 7–17, -3.70 percentage points in low-chance provinces versus -1.42 in high-chance provinces, a between-group difference of -2.28 points (significant at 5%). Panel B adds exposure interactions with four 1990 prefecture characteristics (agricultural employment share, college and secondary attainment stocks, log population), all instrumented: the group difference does not shrink but grows to -2.98 points (1%), while the high-chance provinces' own effect goes to zero ($+0.28$, insignificant). An explanation that attributes the difference to prefecture-level poverty or industrial-structure gradients must first explain why prefecture-level 1990-

characteristic interactions do not absorb it. Figure 7 shows the same contrast age by age: both curves fall steeply after age 15, but the low-chance curve is systematically lower already at compulsory-schooling ages.

Table 7. Stage-by-stage levels in-school effect by pre-expansion (1998) provincial university admission chance, quartile tails. The 22 proper provinces in the sample (the four municipalities carry no valid admission-chance measure and are excluded from the grouping) are split at the top vs. bottom quartile of the 1998 admission chance (provincial regular-college admissions over regular senior-high graduates, 1998 Education Statistical Yearbook; six provinces per tail, middle half excluded). Each column is a separate 2SLS regression for one age window on prefecture × wave × age cells; rows report the top-tail slope, the bottom-tail slope, and their difference, in percentage points per 1 SD exposure level. Panel B adds exposure interactions with four 1990 prefecture characteristics (agricultural employment share, college and secondary attainment stocks, log population), all instrumented. Prefecture, province × wave, and age × group FE; prefecture-clustered SE.

	2SLS: robot exposure (1 SD), pp			
	Primary 7–11 (1)	Junior high 12–14 (2)	Senior high 15–17 (3)	Pooled 7–17 (4)
Panel A: no controls				
Top-quartile provinces	-0.09	-0.74**	-4.32***	-1.42***
	(0.18)	(0.29)	(0.74)	(0.28)
Bottom-quartile provinces	-2.07**	-3.66***	-6.47***	-3.70***
	(0.97)	(1.10)	(1.03)	(0.87)
Bottom–top difference	-1.97**	-2.92**	-2.15*	-2.28**
	(0.99)	(1.14)	(1.27)	(0.92)
First-stage F	651	643	643	658
Panel B: with 1990 characteristics × exposure				
Top-quartile provinces	1.27	1.24	-2.34*	0.28
	(0.85)	(0.94)	(1.22)	(0.73)
Bottom-quartile provinces	-1.42	-2.32**	-5.22***	-2.70***
	(0.99)	(1.04)	(1.09)	(0.84)
Bottom–top difference	-2.69**	-3.56***	-2.88*	-2.98***
	(1.11)	(1.33)	(1.46)	(1.06)
First-stage F	47	47	47	48
Observations	2,460	1,476	1,476	5,412
Fixed effects	Prefecture, province × wave, and age × group FE			
Cell unit	Prefecture × wave × age			
Waves	2000, 2005, 2010, and 2015			
Weights	Equal cells			

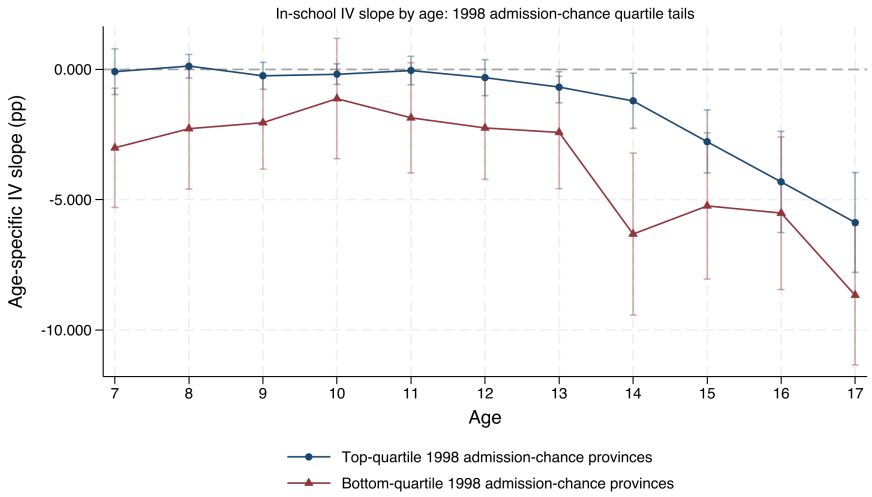


Figure 7. Robot exposure and in-school rate by age, top vs. bottom 1998 admission-chance quartile provinces, in levels. Each group–age slope is estimated in its own regression of the in-school rate on the exposure level with prefecture and province-by-wave fixed effects, instrumented with the EU shifter. Whiskers are 95% CIs.

Three qualifications bound this reading. First, the group difference is most precise at primary and junior-high ages (7–14), while at the senior-high gateway (15–17) both groups show large, highly significant declines and the difference is significant only at 10%—consistent with the option value operating along the whole schooling path (each year enrolled preserves the option of eventually reaching higher education), but "the difference is least precise exactly at the gate" remains a boundary worth recording: it suggests a large share of the gateway decline is a routine-pathway devaluation common to both groups, which the admission-chance difference only modulates. Second, the contrast relies on the quartile tails: under a median split the group differences are same-signed but generally insignificant (not reported here), so the table should be read as a tails contrast rather than a full gradient in the chance rate. Third, as noted in Section 3, admissions are counted at the institution's province, so large sending provinces like Shandong may be over-assigned to the low-chance group. These qualifications match the stance of Section 2: this subsection offers ecological evidence consistent with the option-value channel, not an individual-level test that low-π students leave earlier.

6.4 Household college education: the gateway response concentrates in households without college-educated members

The admission-chance contrast varies at the province level. This subsection brings the same corollary down to the household level, where the progression probability π plausibly varies most: within prefectures, children are split by whether any co-resident household member holds a bachelor-and-above degree (with junior-college-and-above as a variant). The flag is built household by household in all four census waves; collective households are excluded because co-residents there are not family members, and the 2010 household identifiers are rebuilt from the enumeration area $\times$ household serial (the public-use person id does not group households). The share of children living with a bachelor-educated member rises from 1.3% in 2000 to 7.4%

in 2015. The design repeats the levels specification on prefecture–wave–age cells split by the flag, with every fixed effect interacted with the group (prefecture × group, province × wave × group, age × group), so each group's slope is identified within its own prefecture and wave structure and the base coefficient equals the within-group split estimate exactly.

Table 8. Stage-by-stage levels in-school effect by co-resident household college education. Children aged 7–17 in family households across the four census waves are split by whether any co-resident household member holds a bachelor-and-above degree. Each column is a separate 2SLS regression for one age window on prefecture × wave × age × group cells; rows report the slope for households without such a member, the slope for households with one, and their difference, in percentage points per 1 SD exposure level. All fixed effects are interacted with the group; prefecture-clustered SE. Collective households are excluded (co-residents are not family); the clean-school rule is applied at build time.

	2SLS: robot exposure (1 SD), pp			
	Primary 7–11 (1)	Junior high 12–14 (2)	Senior high 15–17 (3)	Pooled 7–17 (4)
Households without a bachelor-educated member	-0.38**	-0.95***	-3.36***	-1.35***
	(0.17)	(0.26)	(0.41)	(0.21)
Households with a bachelor-educated member	-0.34	-0.72*	-0.56	-0.49*
	(0.23)	(0.40)	(0.78)	(0.26)
Difference (with − without)	0.04	0.23	2.79***	0.86***
	(0.28)	(0.43)	(0.81)	(0.29)
First-stage F	258	244	254	259
Observations (cells)	10,944	6,549	6,504	24,008
Fixed effects	Prefecture × group, province × wave × group, age × group			
Cell unit	Prefecture × wave × age × group			
Waves	2000, 2005, 2010, and 2015			
Weights	Equal cells			

Table 8 shows that the two groups respond alike through the compulsory stages—at ages 7–11 and 12–14 the group differences are $+0.04$ and $+0.23$ percentage points, both insignificant—and then split apart exactly at the senior-high gateway: at ages 15–17 the in-school response is -3.36 points (1%) in households without a bachelor-educated member, essentially the paper's headline gateway effect, against -0.56 (insignificant) in households with one, a difference of $+2.79$ points (1%). Pooled over 7–17 the difference is $+0.86$ points (1%). The heterogeneity thus opens precisely at the fee-charging, optional senior-high margin, and it is sharpest exactly where the provincial admission-chance contrast is least precise—the two cuts are complementary: one varies the width of the route upward across provinces, the other the household's own attachment to that route, and both point to the option-value channel.

Three qualifications parallel the previous subsection. First, the flag is a household-composition variable, correlated with urban hukou, income, and location; the fully group-interacted fixed effects give each group its own prefecture and province-wave structure, but the split remains descriptive heterogeneity, not a causal interaction. Second, "any co-resident college-educated member" is not strictly parental education: elder siblings and other adults count, and children whose college-educated parents live elsewhere (including migrant parents) are misclassified into the without-group—a force that, if anything, attenuates the contrast. Third,

unlike the main sample, children in collective households are excluded (about 44,000 in 2010; the coverage file records the cost by wave), and the bachelor group is thin in the early waves.

7 Conclusion

This paper places the question of robots and educational choice at the pre-college stages. Using the four Chinese census and micro-census waves of 2000–2015 and a shift-share instrument built from European sector-level robot diffusion, I find that a one-standard-deviation increase in prefecture robot exposure lowers the in-school rate by -0.42^{**} percentage points at ages 7–11, -1.07^{***} at ages 12–14, and -3.75^{***} at ages 15–17; the age-by-age profile steepens systematically toward the junior-to-senior-high gateway, and the endpoint long difference reproduces the same ordering. This is not a reversal of the schooling expansion: the senior-high-age in-school rate rose from 55.1% to 81.8% over the same period, and the long difference means that a prefecture with a one-standard-deviation larger increase in exposure forwent 6.53 points of that rise. The senior-high slope survives leave-one-sector-out estimates, pre-period characteristic trends, migration restrictions, alternative weighting, and Rotemberg diagnostics, but the most conservative shock-level standard error leaves the estimate imprecise—a qualification that should travel with every reading of the paper.

Two complementary cuts tie the gradient to progression access. At the province level, measured by admission chances on the eve of the 1999 expansion (1998), the school-leaving response concentrates in provinces where the route upward was narrow and is close to zero where it was wide. At the household level, households with and without a co-resident college-educated member respond almost identically through the compulsory ages and then split cleanly at the 15–17 gateway, with the negative response coming almost entirely from households with no member holding a bachelor's degree or above. Rural-hukou teenagers also show more negative point estimates, with the disadvantage again concentrated at the *zhongkao* margin. The first cut varies the width of the route upward across provinces, the second characterizes a family's attachment to that route, and both are consistent with reachable progression opportunities cushioning automation-induced school-leaving. All of this, however, is ecological and descriptive evidence: the provincial contrast relies on quartile tails and is least precise exactly at the gateway, the household-education split measures family composition rather than individual prospects, and the census observes neither test scores, subjective progression probabilities, nor expected wages. These results are therefore offered as channel consistency, not as direct identification of the option-value mechanism.

How this result sits beside the existing evidence is the point the paper most wants to make. The direct evidence from advanced economies is uniformly positive (Dauth et al., 2021; Gunadi and Ryu, 2023; Di Giacomo and Lerch, 2026; Branco et al., 2025), and China too shows positive results for college-age and senior-high-age groups (Liu et al., 2025). The negative sign here does not refute those conclusions: the estimand differs (stopping at age 17, where residence is not yet determined by admission outcomes), as do identification and coverage (a European instrument, a four-wave panel, and an age-by-age profile spanning rural and urban youth). What can be placed alongside them is a general reading: whether automation raises schooling depends on whether the education system can deliver the marginal student to the skill tracks that benefit from the new technology. Where progression is ample, falling opportunity costs and a rising skill premium can be converted into more education; where progression is rationed, continuing in secondary school mainly bets on the routine work automation devalues, and the same shock can push the marginal

student out. The policy implication accordingly lies upstream rather than at the college end: cushioning the educational cost of automation requires acting on capacity and access at the *zhongkao* gateway. That is a directional inference consistent with these estimates, not a welfare evaluation—this paper does not measure the adult consequences of leaving school.

Finally, the boundaries this paper deliberately leaves. The analysis stops at age 17: from age 18 onward, college admission moves residence together with the hukou registration to the prefecture of the admitting institution, prefecture assignment becomes endogenous to the schooling outcome, and the cross-sectional structure of the census cannot track those moves. The extensive margin, moreover, shows only whether a child is still in school, not whether students who remain enrolled increased their exam-preparation effort as the higher-education premium widened. Determining which channel carries the main result—and how large a scar early school-leaving leaves in adulthood—requires linked individual data that can follow people across regions: family income shocks, test scores, and post-18 attainment and labor-market outcomes.

A Appendix: Comparative Statics and the Dropout-Ladder Derivation

This appendix extends the group comparative statics of Section 2 and derives the age-varying dropout ladder; notation follows the text. The in-school share is given by (5); let local robot exposure (denoted x) act through $\partial V^R/\partial x = -\rho$, $\partial V^H/\partial x = -\eta$, and $\partial w/\partial x = -\omega$. The text's exposure-gradient assumption (equation (3)) requires $\rho > 0$ and $\eta < \rho$ and pins the pass-through to the dropout's immediate wage at zero ($\omega = 0$); this appendix keeps $\omega \geq 0$ as a perturbation parameter to characterize the robustness region of that anchoring assumption.

The exposure slope. Differentiating (5) with respect to x,

$$\frac{\partial S_g}{\partial x} \;=\; -f(\cdot)\, L(\pi), \qquad L(\pi) \;\equiv\; \pi\eta + (1-\pi)\rho - \omega,$$

whose sign is determined by the ordering of the three pass-through parameters. Under the text's $\omega = 0$, $L(\pi) = \rho - \pi(\rho - \eta)$ declines linearly in π from $L(0) = \rho > 0$, and the slope is strictly negative on $\pi < \pi^* \equiv \rho/(\rho - \eta)$; when $\eta \geq 0$, $\pi^* \geq 1$ and the slope is negative everywhere. More generally, $\min_{\pi\in[0,1]} L = L(1) = \eta - \omega$, so *the slope is strictly negative for all* $\pi \in [0, 1)$ *if and only if* $0 \leq \omega \leq \eta$: when the college track is devalued too ($\eta \geq 0$), the anchoring assumption can be relaxed to "pass-through to the dropout wage does not exceed pass-through to the college track" without changing the conclusion. Far beyond that bound—for example, one-for-one pass-through of the shock to the dropout's immediate wage ($\omega = \rho$)—$L(\pi) = \pi\,(\eta - \rho) < 0$ for all $\pi > 0$ and the slope turns positive: the opportunity cost of dropping out falls faster than the option value of staying, so enrollment is predicted to rise rather than fall, which is exactly the opportunity-cost logic of the US evidence. An everywhere-negative exposure slope is therefore itself a substantive restriction: robots devalue the terminal value of the "high-school–routine ladder" while passing through almost nothing to dropout teenagers' immediate employment, rather than depressing all low-skill teenage earnings proportionally.

The chance-by-exposure interaction. Differentiating the exposure slope once more with respect to π, the full cross-partial is

$$\frac{\partial^2 S_g}{\partial x\, \partial \pi} \;=\; f(\cdot)\,(\rho - \eta) - f'(\cdot)\left(V^H - V^R\right) L(\pi)$$

(holding ω constant in π). The loading term $f(\cdot)(\rho - \eta) > 0$: the higher is π, the smaller the weight $1 - \pi$ on the most-devalued V^R and the larger the weight on the lightly-hit V^H, so the same robot shock depresses the in-school share less. When F is approximately uniform near the decision margin ($f' \approx 0$), the sign is given by the loading term and the cross-partial is positive; equivalently $\partial\left|\partial S_g/\partial x\right|/\partial\pi = -f(\cdot)(\rho - \eta) < 0$: the larger is π, the smaller the robot effect. Combined with the direction property: on the empirically relevant range $\pi < \pi^*$, the slope is negative with an absolute value decreasing in π. The paper uses this interaction only as a qualitative prediction.

The dropout ladder. Equation (1) places the student at the gate; extending the same decision to younger ages adds two features that vary systematically with age. The first is information. Let the gate be at age A (the school year of the *zhongkao*); a student of age $a < A$ holds only a noisy belief about π: a posterior mean π_a and residual uncertainty σ_a that increases with the years remaining to the gate and vanishes at the gate—the formalization of the text's "π is revealed most sharply at the gate." Leaving school is (nearly) irreversible, and each additional year in school reveals another year of grades, rankings, and admission information, so the value of staying contains an information option O_a: students who stay preserve the right to choose again after the noise resolves. By convexity of the max operator, O_a increases in σ_a—larger the farther from the gate—and converges to zero at the gate.[4] The second is the legal working age. The Labor Law prohibits employing those under 16, a threshold that falls right around the year of the senior-high sorting: younger dropouts cannot access w and are left with informal or within-household work. Write the exit value as $\kappa_a\, w$, with employment accessibility $\kappa_a < 1$ rising in age and approaching one only near the gate. The age-a in-school share then generalizes to

$$S_a \;=\; F\big(\pi_a V^H + (1-\pi_a)\, V^R + O_a(\sigma_a) - \kappa_a\, w - c_a\big), \tag{10}$$

and the exposure slope becomes (keeping the text's $\omega = 0$, so the dropout-wage term carries no loading)

$$\frac{\partial S_a}{\partial x} \;=\; -f(N_a)\,\big[\pi_a\eta + (1-\pi_a)\rho\big], \qquad N_a \equiv \pi_a V^H + (1-\pi_a)V^R + O_a - \kappa_a w - c_a.$$

The loading $\pi_a\eta + (1-\pi_a)\rho$ is linear in beliefs, so compressing the posterior mean does not change its expectation: the expected devaluation enters younger students' staying value at full size, and the loading itself does not vary systematically with age. The ladder is therefore carried entirely by the marginal-density multiplier: at young ages O_a is large and κ_a small, the net value N_a sits deep on the staying side, the decision falls in the thin tail of F, $f(N_a)$ is small, and the same value shock moves few students; approaching the gate, the information option is exhausted and $\kappa_a \to 1$, N_a moves into the thick part of the density, and the mass of students standing on the margin is largest. The modeled dropout ladder—the negative slope steepening step by step as age approaches the senior-high gateway—is thus a shape prediction about the marginal density thickening with age, exactly the shape tested by the age-by-age profile of Figure 4 and the stage regressions. Young ages are not thereby a placebo: the expected V^R devaluation still enters younger students' staying value, but it is pushed into the thin tail of the distribution by the information option and the legal working age.

Cross-region implication. The level effect of progression chances and the chance-by-exposure interaction jointly give the cross-region implication: where π—the probability of continuing through secondary school to higher education—is higher, the in-school share is higher in levels and the absolute exposure slope is smaller; and under the exposure-gradient assumption the slope does not change sign, so the prediction for high-chance regions is a negative effect decaying toward zero, not a positive one. Section 6.3 tests this

[4] Formally this is an optimal-stopping problem: the continuation value takes the form $E_a[\max\{\cdot\}]$, which exceeds $\max\{E_a[\cdot]\}$, and the difference is O_a. Equation (10) is a reduced-form shorthand, not a full dynamic-programming solution.

implication with provincial university admission chances on the eve of the expansion (1998): the quartile-tails contrast shows the exposure response is indeed concentrated in low-chance provinces, and the high-chance response goes to zero rather than positive once prefecture-characteristic interactions are added; but the difference is insignificant under a median split and weakest at the senior-high gateway, so the paper reads it as ecological evidence consistent with the option-value channel, without treating "high-π students stay" as established at the individual level.

B Appendix: Robustness Tables

This appendix collects the robustness tables summarized in Section 5.

Table A1. Weighting-scheme sensitivity of the pooled stage-by-stage levels in-school slope. Each cell is the pooled 2SLS levels in-school slope for one schooling stage under one weighting scheme (equal cells, raw population counts, within-year normalized counts). All specifications include prefecture, province-by-wave, and age FE across the four waves. Negative coefficients mean lower enrollment in more-exposed prefectures; prefecture-clustered SE in parentheses.

	Equal cells (1)	Raw counts (2)	Within-year (3)
Primary 7–11	-0.42**	-0.52***	-0.37***
	(0.17)	(0.14)	(0.13)
Junior high 12–14	-1.07***	-1.54***	-0.91***
	(0.24)	(0.28)	(0.19)
Senior high 15–17	-3.75***	-7.05***	-3.01***
	(0.39)	(0.70)	(0.51)
Estimator	2SLS pooled levels in-school slope; pp per 1 SD exposure level		

Table A2. Migration-restriction robustness by schooling stage, in levels, pooled across hukou. The sample keeps youth whose five-year-prior residence was local or comparable within province, across the four census waves. All specifications include prefecture, province × wave, and age FE. Negative coefficients mean lower enrollment in more-exposed prefectures. Coefficients are percentage points per 1 SD exposure level. Windows are 7–11, 12–14, and 15–17. 2SLS; prefecture-clustered SE; equal-cell weights.

	Primary 7–11 (1)	Junior high 12–14 (2)	Senior high 15–17 (3)
Migration-restricted	-0.05*	-0.76***	-3.62***
	(0.03)	(0.15)	(0.36)
Observations	5,398	3,238	3,239
First-stage F	505.9	500.5	500.5
Fixed effects	Prefecture, province × wave, and age FE		
Cell unit	Prefecture × wave × age (pooled across hukou)		
Estimator	2SLS levels; coefficients in pp; negative means lower enrollment		
Weights	Equal cells		

Table A3. Rotemberg weights for the EU-adoption Bartik instrument in the pooled five-year-difference senior-high in-school effect, ages 15–17, residualized on province-by-period FE. α_k is the sector sensitivity weight and $\hat{\beta}_k$ is the just-identified IV estimate. This table and Table A5 use the same sample, window, standardization, and fixed effects; Table 4 instead uses the 2000–2015 endpoint change.

Sector	Rotemberg weight α_k	2000 emp. share (%)	EU shock change ending 2015	$\hat{\beta}_k$ (pp)
Automotive and other vehicles	0.388	0.69	11.1	-3.83
Electronics	0.186	1.00	1.5	-3.73
Industrial machinery	0.147	0.99	4.0	0.90
Plastics and chemicals	0.090	1.26	1.2	-6.32
Textiles	0.078	2.45	0.8	-1.91
Sum of top-5 weights		0.887		
Sum of negative weights (1 sectors)		-0.039		
Effective number of sectors ($1/\sum_k \alpha_k^2$)		4.4		
IV estimate $\sum_k \alpha_k \hat{\beta}_k$ (pp)		-2.91		
Weights		Equal cells		

Table A4. Leave-one-sector-out robustness for the pooled senior-high-age in-school slope, ages 15–17. Each specification re-estimates the senior-high in-school slope (Table 3) after dropping one IFR sector from both the China exposure and the European shifter. The table reports the all-sector estimate, two high-weight exclusions, and the minimum/maximum leave-one-out estimates; the full sector-by-sector output is in `output/tables/t_leave_one_sector.csv`. Negative coefficients mean lower senior-high-age in-school rates. 2SLS; prefecture-clustered SE; equal-cell weights.

Specification	Coef. (pp)	SE	First-stage F
All sectors	-3.75***	(0.39)	520.7
Automotive and other vehicles	-3.06***	(0.46)	173.8
Electronics	-3.82***	(0.53)	1747.1
Leave-one-out range	-3.85 to -3.06	–	–
Range endpoints: Industrial machinery and Automotive and other vehicles.			
Negative leave-one-out estimates	18 of 18 sectors		
Fixed effects	Prefecture, province × year, and age FE		
Weights	Equal cells		

Table A5. Shock-level inference for the pooled long-difference senior-high-window in-school effect, ages 15–17, residualized on province-by-period FE. Shock-level SEs follow Borusyak et al. (2022).

	In school
Senior-high-window exposure effect (pp)	-2.91
Prefecture-clustered SE (268 clusters)	(0.34)***
Shock-level robust SE (53 shocks)	(1.93)
Shock-level sector-clustered SE (18 sectors)	(0.83)***
Prefecture-period cells	796
Sector × wave shocks (nonzero)	53
Weights	Equal cells

References

Abramitzky, R. and V. Lavy (2014). How Responsive Is Investment in Schooling to Changes in Redistributive Policies and in Returns? *Econometrica* 82(4), 1241–1272.

Acemoglu, D. and P. Restrepo (2018). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review* 108(6), 1488–1542.

Acemoglu, D. and P. Restrepo (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy* 128(6), 2188–2244.

Acemoglu, D. and P. Restrepo (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica* 90(5), 1973–2016.

Adão, R., M. Kolesár, and E. Morales (2019). Shift-Share Designs: Theory and Inference. *Quarterly Journal of Economics* 134(4), 1949–2010.

Adukia, A., S. Asher, and P. Novosad (2020). Educational Investment Responses to Economic Opportunity: Evidence from Indian Road Construction. *American Economic Journal: Applied Economics* 12(1), 348–376.

Atkin, D. (2016). Endogenous Skill Acquisition and Export Manufacturing in Mexico. *American Economic Review* 106(8), 2046–2085.

Autor, D. H., F. Levy, and R. J. Murnane (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *Quarterly Journal of Economics* 118(4), 1279–1333.

Borusyak, K., P. Hull, and X. Jaravel (2022). Quasi-Experimental Shift-Share Research Designs. *Review of Economic Studies* 89(1), 181–213.

Branco, D., B. Carrillo, and W. Iglesias (2025). Routine-Biased Technological Change and Endogenous Skill Investments. *American Economic Journal: Economic Policy* 17(3), 236–280.

Charles, K. K., E. Hurst, and M. J. Notowidigdo (2018). Housing Booms and Busts, Labor Market Opportunities, and College Attendance. *American Economic Review* 108(10), 2947–2994.

Dauth, W., S. Findeisen, J. Südekum, and N. Wößner (2021). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association* 19(6), 3104–3153.

Di Giacomo, G. and B. Lerch (2026). Automation and Human Capital Adjustment: The Effect of Robots on College Enrollment. *Journal of Human Resources* 61(3), 733–760.

Giuntella, O., Y. Lu, and T. Wang (2025). How Do Workers Adjust to Robots? Evidence from China. *Economic Journal* 135(666), 637–652.

Goldin, C. and L. F. Katz (2008). *The Race between Education and Technology*. Cambridge, MA: Belknap Press of Harvard University Press.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020). Bartik Instruments: What, When, Why, and How. *American Economic Review* 110(8), 2586–2624.

Greenland, A. and J. Lopresti (2016). Import Exposure and Human Capital Adjustment: Evidence from the U.S. *Journal of International Economics* 100, 50–60.

Gunadi, C. and H. Ryu (2023). How Do People Respond When They Know That Robots Will Take Their Jobs? *Oxford Bulletin of Economics and Statistics* 85(5), 939–958.

Hannum, E. (1999). Political Change and the Urban–Rural Gap in Basic Education in China, 1949–1990. *Comparative Education Review* 43(2), 193–211.

Jensen, R. (2010). The (Perceived) Returns to Education and the Demand for Schooling. *Quarterly Journal of Economics* 125(2), 515–548.

Kang, R. (2026). The Dynamic Effects of Robot Exposure on Human Capital Investment and Labor Market Outcomes. Job market paper. https://sites.google.com/view/rihyeon-kang/research

Li, B. (2018). Export Expansion, Skill Acquisition and Industry Specialization: Evidence from China. *Journal of International Economics* 114, 346–361.

Li, H., P. W. Liu, and J. Zhang (2012). Estimating Returns to Education Using Twins in Urban China. *Journal of Development Economics* 97(2), 494–504.

Liu, Y., Y. Guo, Z. Li, B. Wang, and J. You (2025). The Human Capital Promotion Effect of Industrial Intelligence: Theory and Evidence. Working paper, SSRN No. 5553525. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5553525

Loyalka, P., C. Liu, Y. Song, H. Yi, X. Huang, J. Wei, L. Zhang, Y. Shi, J. Chu, and S. Rozelle (2013). Can Information and Counseling Help Students from Poor Rural Areas Go to High School? Evidence from China. *Journal of Comparative Economics* 41(4), 1012–1025.

Stange, K. M. (2012). An Empirical Investigation of the Option Value of College Enrollment. *American Economic Journal: Applied Economics* 4(1), 49–84.

World Bank (2026). School Enrollment, Tertiary (% Gross), Indicator SE.TER.ENRR. *World Development Indicators*. https://data.worldbank.org/indicator/SE.TER.ENRR?end=2024&locations=US-CN&start=1970&view=chart, data version 2026-07-13, accessed 2026-08-18.

Yi, H., Y. Song, C. Liu, X. Huang, L. Zhang, Y. Bai, B. Ren, Y. Shi, P. Loyalka, J. Chu, and S. Rozelle (2015). Giving Kids a Head Start: The Impact and Mechanisms of Early Commitment of Financial Aid on Poor Students in Rural China. *Journal of Development Economics* 113, 1–15.

Zhang, J. and K. Zhou (2023). Quota Removal, Destination-Specific Export Shocks, and Skill Acquisition in China. *Journal of Development Economics* 165, 103149.

Zhang, R. (2026). Higher Education Expansion and the Rural–Urban Gap in Secondary Schooling and Labor Outcomes: Evidence from China's 1999 Reform. *Journal of Development Economics* 181, 103721.