

The Impact of Trade Liberalization on School Enrolment: Evidence from China

Hanqi Liu*

University of Calgary

[Click here for the latest version](#)

October 16, 2026

Abstract

This paper exploits quasi-experimental variation in local exposure to the shock of China's accession to the World Trade Organization (WTO) to evaluate its impact on high school enrolment, a key margin along which human capital accumulates. I construct a measure of prefecture-level trade exposure using initial industrial composition, industry-level tariff reductions, and spatially determined geographic trade costs, and I estimate both the short- and long-run impact of the shock on enrolment. Using detailed population census data on successive cohorts of 16-year-olds from 1995 to 2014, I find that a one-standard-deviation increase in exposure reduces high school enrolment by 8% relative to the baseline rate of the last pre-accession cohort. Surprisingly, this drop in enrolment persists and even grows across successive cohorts. These results suggest that China's WTO accession created an export-driven labour demand shock biased toward low-skilled workers and reinforced by subsequent complementary investment.

Keywords: Trade liberalization, school enrolment, trade costs, low-skilled workers

JEL Codes: F16, F66, J23, J24

*I thank M. Scott Taylor for his continuous guidance and support, and editor Steven Lehrer and two anonymous referees for their very helpful comments. I am also grateful to Jean-William Laliberté, Arvind Magesan, Jiatong Zhong, and Brian McCaig for their comments on this, my second-year paper. Finally, I thank conference participants at the 2024 Canadian Economic Association Meetings (Toronto) and the Inaugural Meeting of the Canadian International Trade Study Group (Burlington) for their valuable input.

1 Introduction

Since the 1980s, many developing countries have undergone significant trade liberalizations. The resulting export expansion has created millions of jobs and increased household incomes. The hope is that these gains will also improve the human capital of the next generation: as household incomes rise, families can better afford schooling and children stay in school longer. While this hope is intuitive, the causal effect of trade liberalization on schooling is not fully understood, especially over the longer run. In this paper, I investigate the impact of the China Shock—China’s accession to the World Trade Organization (WTO)—on high school enrolment within China and provide new evidence on how this effect varies across cohorts.

To explore the underlying mechanisms, I present a simple conceptual framework where family income and relative wages jointly affect the schooling decision. In China, free compulsory education ends after middle school, and families must pay tuition and fees to send a child to high school.¹ If the trade shock increases family income, it eases this financial burden and encourages families to send their children to high school, an income effect. But the same export boom can also expand manufacturing industries that rely heavily on low-skilled workers, raising the demand for these workers and their wages. Higher wages for low-skilled workers can encourage youths to forgo further schooling for immediate employment, a substitution effect. Which effect dominates is an empirical question.

A further question is how the shock’s impact might change over time. Does the one-time reallocation of labour towards comparative-advantage sectors create only a temporary change in the incentives for further schooling, one that dissipates as the economy adjusts? Or do these changes instead intensify over time, as firms install capital complementary to low-skilled labour and further raise demand for these workers? These two possibilities imply different patterns across cohorts ranging from those making their schooling decisions just after the shock to those deciding more than a decade later. Comparing cohorts therefore allows me to distinguish between these possibilities.

To shed light on these questions, I exploit quasi-experimental variation created by China’s WTO accession in December 2001. This accession has two features that make it well suited for identification. First, it was plausibly exogenous to household schooling decisions. The timing of accession was difficult to anticipate, as the process involved prolonged and uncertain bilateral and multilateral negotiations driven by China’s complex international rela-

¹In China, middle school covers grades 7 to 9, and high school covers grades 10 to 12. Students typically complete middle school at age 15 or 16. High school costs can be a substantial burden, especially for rural families; see Section 2.1 for detailed background on China’s secondary schooling system.

tions, which left families little scope to adjust their schooling decisions in advance. Second, the resulting export expansion was unusually large. After joining the WTO, China's gross exports surged from \$0.27 trillion in 2001 to \$2.34 trillion in 2014, making it the world's largest goods exporter. Employment in manufacturing rose alongside this expansion, from 100 million workers in 2002 to 150 million in 2008. Manufacturing draws heavily on less-skilled workers: in 2000, over 72% of working-age employees in manufacturing had less than a high school degree. The expansion of this sector disproportionately raised demand for workers without a high school education.²

To investigate how the trade shock affected high school enrolment, I use population census data covering successive cohorts turning 16 between 1995 and 2014. I focus on age 16 because it is when students typically finalize their high school entry decisions. Enrolment records whether students go on to high school, and I use it as a proxy for education.³ I construct a measure of prefecture-specific trade exposure by combining initial industrial composition, industry-level tariff reductions, and spatially determined geographic trade costs (hereafter, "trade costs").⁴ A prefecture's trade exposure depends jointly on its industrial composition, which governs the potential shock from the tariff reductions, and its trade costs, which determine how strongly this potential translates into an actual trade shock. I compare cohorts who made the high school schooling decision before and after the shock across prefectures that differ in their levels of trade exposure.

I find that greater exposure to China's WTO accession reduces high school enrolment among 16-year-olds. Prefectures with one standard deviation higher exposure see high school enrolment fall by an additional 3.46 percentage points relative to less exposed ones, an 8% decline relative to the baseline enrolment rate of the cohort turning 16 in 2001, the last pre-accession cohort. While trade liberalization can theoretically induce competing income and substitution effects, this reduced-form result indicates that the substitution effect dominates. To unpack this aggregate effect, I separate enrolment by track. In China, students who continue beyond middle school enter either academic schools (preparing for university) or vocational schools (leading directly to the labour market). I find that the overall decline is driven almost entirely by the academic track. In contrast, the estimated effect on vocational enrolment is positive, albeit statistically insignificant. To trace how

²The export data are from the Trade Analysis Information System database ([TRAINS](#)); the manufacturing employment data are from the Socio Economic Accounts (2014) from the [World Input-Output Database](#); the educational attainment of manufacturing employment is calculated from the 2000 Population Census.

³Education is multidimensional ([Hanushek and Woessmann, 2008](#)), and I focus on enrolment, its quantity side, which can be measured consistently across prefectures and over time in China, whereas its quality side varies sharply both within and across prefectures and is difficult to measure; see Section 2.1 for details.

⁴In China, a prefecture is an administrative unit between the province and the county.

the effect evolves across cohorts, I employ an event-study model allowing for year-by-year variation in the estimated effects. The negative effect of trade exposure on high school enrolment not only persists but grows across cohorts through the end of the sample period in 2014.

Why does the effect of trade on schooling intensify rather than fade? I present descriptive evidence on two forces that together plausibly reinforce the initial shock, one operating through labour demand and the other through China's schooling system. In provinces with greater exposure, capital investment surged following the shock, expanding productive capacity in labour-intensive export sectors.⁵ Because this capital is complementary to low-skilled workers, its continued installation is consistent with demand for such workers persisting well beyond the shock. Facing this sustained demand, students who might otherwise have entered the academic track appear to have moved into the workforce or, to a smaller and statistically insignificant degree, into the vocational track. Because schooling decisions made around age 16 are hard to reverse under China's schooling system, each cohort that leaves the academic track adds durably to the supply of low-skilled workers, so these two forces compound across cohorts. The educational distribution shifts accordingly: a smaller share of each cohort attains tertiary education and a larger share stops at the compulsory level. Rather than a broad-based upgrading of educational attainment, the estimates show a shift in the distribution of attainment across successive cohorts toward the lower end.

The persistence of the negative effect on high school enrolment mirrors the long-lived labour market impacts of the China Shock in the U.S., where persistence works through local labour markets that adjust gradually (Autor, Dorn and Hanson, 2021). In China, it works through the schooling decisions of successive cohorts, which are difficult to reverse once made. As global competition intensifies and automated technologies increasingly substitute for routine labour (Acemoglu and Restrepo, 2022), cohorts that left school earlier may find it harder to move out of routine work, and high-exposure regions may struggle to transition away from labour-intensive manufacturing.

This paper differs from the existing literature on China's trade liberalization and schooling in three respects: it incorporates trade costs into the measure of trade exposure, it distinguishes the immediate effect on cohorts deciding around the accession from the longer-run effects on subsequent cohorts, and it examines the mechanism behind the persistence of these effects and their consequences for schooling tracks and the distribution of attain-

⁵Given data availability, I present the descriptive evidence at the province level instead of the prefecture level.

ment.

A first difference lies in the treatment of geographic heterogeneity. While studies such as Li (2018), Li et al. (2019) and Jiang, Kennedy and Zhong (2023) use industrial composition to capture the unequal effect of trade shocks across prefectures, these earlier works do not account for differences across prefectures in their spatially determined geographic trade costs. This omission matters because I find that trade costs vary substantially across prefectures. In general, eastern coastal prefectures have lower trade costs than inland prefectures, but trade costs do not simply increase with distance to the coast, because they also reflect differences in infrastructure and policy support.⁶

Omitting trade costs introduces omitted-variable bias, the size of which highlights the importance of geography in constructing the exposure measure. I find that trade costs account for 75.3% of the cross-prefecture variation in trade exposure, compared to 58.9% for the industry-side trade shock (industrial composition weighted by tariff reductions). Moreover, prefectures with lower trade costs tended to face a larger industry-side trade shock, which likely reflects their more advanced industrialization around 2000, so these two components are positively correlated.⁷ Among prefectures with the same industry-side trade shock, those with lower trade costs experience larger trade exposure and a larger decline in enrolment. Given this correlation, an exposure measure built on the industry side alone attributes the declines driven by trade costs to the industry-side shock and therefore overstates the magnitude of its negative effect.

A second difference is that I examine how the trade shock persistently reshapes the schooling decisions of successive cohorts. Existing work speaks to this question only in part. First, studies analyzing the 2001 WTO accession examine long-term outcomes with a focus on industry skill specialization rather than the underlying schooling decisions (Li, 2018; Li et al., 2019; Zhang and Zhou, 2023). Second, while Jiang, Kennedy and Zhong (2023) track cohort-level responses, they study the 1978 Open-Door Policy, a spatially concentrated composite reform in a planned agrarian economy without compulsory education. The 2001 accession was instead a nationwide trade shock to an economy with an already substantial manufacturing base and higher rates of middle school completion.

A third difference is that I examine a mechanism and a set of structural consequences

⁶See Figure 3(a) for an illustration of trade costs across Chinese prefectures. Measurement details are in Section 3.2. My estimates are consistent with Li, Yang and Huang (2022) and Tombe and Zhu (2019), who use the Head-Ries Index to estimate provincial trade costs in 2000 and regional trade costs in 2002, respectively.

⁷Trade costs enter the exposure measure in inverse form, so that prefectures facing lower trade costs are more exposed to the tariff-induced export boom.

that the literature has largely overlooked. First, on the mechanism behind the persistent enrolment decline, I present descriptive evidence that continued investment in labour-intensive export sectors is consistent with demand for low-skilled workers remaining elevated across successive cohorts. Second, on the structural consequences, I use a decomposition of the enrolment decline to show that it is driven by a sharp contraction in the academic track, with little offsetting substitution toward the vocational track. Constructing cohort-specific measures of educational inequality, I further show that this contraction leaves overall inequality largely unchanged while shifting the composition of attainment: fewer students reach tertiary education and more stop at the compulsory level.

My findings also contribute to the rich literature on trade liberalization and human capital in developing countries. Two key results are that trade liberalization can have large effects on labour market choices and outcomes (see [Edmonds and Pavcnik \(2005\)](#); [Edmonds, Pavcnik and Topalova \(2010\)](#)), and that the direction of its impact on schooling depends on which industries it favours ([Atkin, 2016](#); [Blanchard and Olney, 2017](#)). My results echo these findings in a very different setting: China was a much larger, more diverse, and relatively closed economy prior to accession.

The structure of this paper is as follows. Section 2 describes the institutional background of secondary schooling in China and the data. Section 3 outlines the empirical strategy. Section 4 presents the main results and robustness checks. Section 5 concludes.

2 Background and data

2.1 Secondary schooling in China

China enacted its Compulsory Education Law in 1986, establishing nine years of compulsory schooling: six years of primary school (Grades 1 to 6) and three years of middle school (Grades 7 to 9, junior middle school). The law requires children to enrol at age six, or at seven where local conditions do not permit earlier entry, and to complete the full nine years. Compulsory education is free of tuition, and the law directs local governments to ensure that school-age children enrol in schools near their place of household registration (*hukou*). Implementation nevertheless remained incomplete well into the 2000s: by the end of 2000, regions accounting for only 85% of the national population had achieved universal nine-year compulsory education ([Ministry of Education of the People's Republic of](#)

China, 2001), with attainment lagging especially in poorer rural areas.⁸

Students normally complete compulsory education at the age of 15 or 16, at which point they face two options. First, they may leave formal schooling; those who have reached the legal working age of 16 can enter the labour market directly. Second, those who wish to continue to high school (Grades 10 to 12, senior middle school) must sit the high school entrance examination (*zhongkao*), which is administered at the prefecture level at the end of middle school.⁹ This examination is the main channel through which students are sorted into two distinct tracks: academic and vocational. Both tracks last three years but serve different purposes. Academic high schools are oriented toward further study, with graduates competing through the National College Entrance Examination (*gaokao*) for places in higher education. Vocational high schools (broadly defined to include specialized secondary schools and skilled-worker schools) are oriented toward employment and train workers for specific industries such as nursing and automobile repair. Admission to both tracks is based on *zhongkao* scores, but the score thresholds differ sharply: academic high schools are selective and admit only students above competitive cutoffs, whereas vocational high schools admit students with much lower scores. Because students clearing the academic cutoff may still choose the vocational track or forgo high school, placement reflects both scores and household choice.

According to the Educational Statistics Yearbook of China, in 2001, academic high schools enrolled 14.1 million students, or 59% of total high school enrolment, while vocational schools enrolled 9.8 million, or 41%. Unlike North American comprehensive high schools, where vocational coursework still leaves the university route open, China's vocational track is institutionally separate and prepares students to enter the workforce upon graduation. During the sample period, the vocational track was essentially terminal, with pathways into higher education remaining marginal until well after.¹⁰ In the years under study, therefore, the academic track was the principal route to university, and exiting it at the *zhongkao* stage effectively foreclosed subsequent access to higher education.

High school, unlike compulsory education, is not free, and the financial burden on families is substantial, especially for rural families sending a child to academic high school. Surveying eight poor rural counties in Shaanxi Province, Liu et al. (2009) find that three

⁸See Zhang and Zhao (2006) for detailed statistics; for example, the population coverage rate for compulsory education stood at only 77% in rural western regions as of 2003.

⁹See Ding and Lehrer (2007) for an account of the *zhongkao* and admission procedures in urban China.

¹⁰While vocational students were not formally barred from the *gaokao*, their curricula focused on occupational skills rather than academic subjects. A separate track-specific examination admitted a small fraction of vocational graduates and led primarily to two- to three-year colleges rather than four-year universities.

years of tuition alone equaled 82% of one year’s rural per capita net income, and non-tuition expenses such as boarding, food, and tutoring added more than double that amount.

Finally, the quality of both middle and high school varies sharply within and across prefectures, and the urban-rural gap is pronounced (Ding and Lehrer, 2007; Lehrer and Wang, 2026; Li et al., 2017). Some families send their children to better-resourced schools outside their prefecture of registration. The *hukou* system, however, keeps these students tied to their place of registration. The *zhongkao* and the *gaokao* must in principle be taken in a student’s *hukou* locality, where admission quotas are allocated, so students who study elsewhere must return to sit these examinations. Access to high school is largely determined by a student’s prefecture of registration and access to university by their province, regardless of where they attend school.¹¹ Given this variation in school quality and the difficulty of measuring it consistently, I use enrolment as a proxy for education.

2.2 Data

Population and educational attainment. The demographic and education data come from five waves of China’s Population Census and intercensal surveys: the 1990, 2000, and 2010 censuses and the 2005 and 2015 One-Percent Population Sample Surveys (mini-censuses).¹² The 1990 and 2000 censuses are obtained from the Integrated Public Use Microdata Series (IPUMS)-International; the remaining waves are obtained from the National Bureau of Statistics of China. These data contain information on age, gender, prefecture of residence, educational attainment, enrolment status, employment industry, and migration status. To make prefectures comparable across waves, I concord prefecture codes to the 2005 administrative classification. I use the 2015 mini-census to measure whether each individual ever attended high school and focus on individuals born between 1979 and 1998, who made high school entry decisions between 1995 and 2014. The census records the highest level of schooling an individual has attended along with enrolment status. I use only the level and do not distinguish between completion, dropout, and current enrolment status. Each individual is assigned to the prefecture of residence reported in the census.¹³

Figure 1 draws on all five census waves to provide an overview of total high school

¹¹Post-2012 reforms permitted some migrant students to take the *zhongkao* and *gaokao* outside their *hukou* locality, subject to stringent parental employment and social-insurance requirements. These changes occurred only toward the end of the sample period.

¹²For 2015, I use the 20% subsample of the China 1% Population Sampling Survey.

¹³See Section 4.6.4 for discussions on migration issues.

enrolment in China, covering both academic and vocational tracks. The enrolment rate is calculated at the prefecture level as the share of individuals aged 16–18—the typical ages for high school attendance—who have ever attended high school. Panels (a)–(e) map the enrolment rate across prefectures for each of the five census years, with darker shades of blue indicating higher enrolment rates. Panel (f) plots the distribution of enrolment rates across prefectures in each year, with dashed lines marking the corresponding means.

Three patterns stand out. First, high school enrolment rose substantially over time. Panel (f) shows that the mean enrolment rate across prefectures increased from 13% in 1990 to 70% in 2015. Second, this expansion was uneven across space. As the maps in Panels (a)–(e) show, within each year enrolment was higher in economically developed areas such as coastal prefectures. The timing of expansion also differed: enrolment rose first in the northeast and along the coast, while central prefectures caught up rapidly after 2005 and southwestern prefectures more slowly, so that enrolment in parts of the southwest remained lower even in 2015.¹⁴ Third, because enrolment expanded at different speeds across prefectures, the distribution of enrolment rates became more dispersed between 2000 and 2010, as Panel (f) shows.

Trade policy and geographic frictions. Trade and tariff data are from the Trade Analysis Information System Database (TRAINS). For each 4-digit International Standard Industrial Classification (ISIC) Rev. 3 industry, I obtain China’s export value and the tariffs imposed by its trading partners. I concord the ISIC export and tariff data to the Chinese Standard Industrial Classification (CSIC, 1984 version), which is the classification system used in the 1990 Population Census.

The trade cost measure draws on three additional sources. First, country-level exports and GDP are from World Bank Open Data. Second, provincial exports and GDP are from the Statistical Yearbooks of the National Bureau of Statistics; all monetary series are converted to 2015 constant U.S. dollars using RMB–USD exchange rates and U.S. CPI deflators from Federal Reserve Economic Data (FRED). Third, I identify major seaports listed in the official National Coastal Port Layout Plan and verify the year in which each began operation, retaining the 38 fully operational by 1990. Their locations predate the WTO accession and cannot have been chosen in response to the subsequent export surge. Using the 2005 prefecture boundaries, I then compute the great-circle distance from each prefecture’s ge-

¹⁴The later catch-up partly reflects national policies such as the 2006 reform strengthening funding for rural compulsory education in less-developed regions, which plausibly raised middle school completion and thereby expanded the pool of students eligible to continue to high school.




Figure 1: High School Enrolment Rate Across Prefectures, 1990–2015.

Notes: Each map plots the high school enrolment rate at the prefecture level, defined as the share of individuals aged 16 to 18 who have ever attended high school of either academic or vocational type. Data are from the 1990, 2000, and 2010 censuses and the 2005 and 2015 One-Percent Population Sample Surveys; individual sampling weights are applied for the latter two. Darker shades indicate higher enrolment rates, and gray areas denote prefectures with missing data. Panel (f) plots the density of enrolment rates across prefectures in each census year, with dashed lines marking the mean across prefectures. This figure covers the 335 prefectures with consistent definitions across all five census waves.

ographic centroid to the nearest of these ports. Appendix B.2 lists and maps these ports.

3 Empirical strategy

This paper investigates the impact of trade liberalization on high school enrolment in China. To isolate this impact, I use four sources of variation: the quasi-experiment created by China’s accession to the WTO, industry-specific tariff reductions, industrial composition, and spatially determined geographic trade costs.

In terms of temporal variation, the post-accession cohorts are teenagers who turned 16 after 2001. For regional variation, I employ a continuous treatment approach that combines

the other three sources of variation above. Prefectures with higher trade exposure are those that combine lower trade costs with industrial compositions more sensitive to tariff reductions. I compare high school enrolment between cohorts who turned 16 before and after the China Shock, across prefectures with varying levels of exposure.

3.1 Four sources of variation

Temporal variation. I use China's accession to the WTO as a quasi-experiment to study how trade liberalization affected school enrolment. The accession was surrounded by uncertainty, so its timing and terms were difficult for households to anticipate when making schooling decisions. This uncertainty came from two sources. First, the negotiations stretched from 1986 to 2001 and were repeatedly interrupted amid a complicated international environment, most notably following the 1989 Tian'anmen Square protests and the 1999 U.S. bombing of the Chinese Embassy in Yugoslavia. Neither the date nor the terms of entry could be predicted in advance. Second, Most Favoured Nation (MFN) status was itself uncertain. Under the MFN principle, any concession granted to one member, such as a lower tariff, must be extended to all. Although China had received MFN tariff rates from the United States since the 1980s and from the European Union since the 1990s, this treatment remained uncertain for exporters, given annual renewal requirements and the threat of antidumping measures.¹⁵

After China joined the WTO, MFN treatment became a standing obligation among members, removing the year-to-year renewal uncertainty and providing exporters with a stable environment for long-term investment and expansion. As the red line in Figure 2 shows, China's manufacturing exports to the global market rose sharply after 2001.

Industry-specific tariff reductions. The stabilization of MFN status substantially reduced the import tariffs faced by Chinese exporters in the manufacturing sector. Between 1990 and 2005, the weighted average tariff in the manufacturing sector declined by over 60%, as shown by the blue line in Figure 2. These reductions also varied across manufacturing industries. Between 2001 and 2014, manufacturing tariffs fell by 5% on average in relative terms, and this change varied across industries, with a standard deviation of 31%. Under MFN principles, tariff rates are set through multilateral negotiations and applied uniformly to all members, so they do not respond to conditions in any individual Chinese prefecture. Industry-level differences in tariff reductions are therefore plausibly exoge-

¹⁵See [Pierce and Schott \(2016\)](#) for more details about uncertainty and its consequences.




Figure 2: Exports and Tariffs in Manufacturing

Notes: The export and tariff data are from the TRAINS database. The dashed vertical line marks 2002, the first post-accession year. The red line plots China's manufacturing exports to the world, calculated using imports from China reported by partner economies, in billions of 2015 constant USD. The blue line plots the average tariff faced by Chinese manufacturing exports, computed as the effectively applied (AHS) tariff levied by each destination on imports from China, weighted by China's exports to each destination in the corresponding years.

nous to local schooling decisions, and provide a second source of identifying variation to analyse the impact of trade liberalization.

Prefecture-specific industrial composition. The impact of industry-specific tariff reductions varies across Chinese prefectures due to differences in their industrial composition before the trade shock. In 1990, prefectures differed substantially not only in the overall proportion of their workforce employed in manufacturing—with an average of 9.6% and a standard deviation of 9.2%—but also in the specific mix of manufacturing industries. This local variation determines how strongly each prefecture is affected by tariff reductions in specific industries. Prefectures with a higher concentration of industries experiencing larger tariff reductions are more exposed to the China Shock.

Spatially determined geographic trade costs. While industry-specific tariff reductions provide useful variation, tariffs are only one component of the broader frictions that shape trade flows. Following [Anderson and Van Wincoop \(2004\)](#), trade costs encompass not only tariffs but also geographic barriers such as transport (freight and time), regulatory and legal compliance, language, and local distribution (wholesale and retail). These non-tariff components are large: for a representative high-income country, the total trade cost is estimated at an ad-valorem tax equivalent of roughly 170%, of which tariffs and other

formal trade-policy instruments account for only a small share, with average tariffs on the order of 5%.

Spatially determined trade costs vary markedly across Chinese prefectures because of both nationwide and locally specific factors. Access to seaports lowers trade costs by enabling cheaper and faster connections to global markets, and only about 17% of Chinese prefectures are coastal. More broadly, prefectures with better infrastructure, including roads, rail, airports, and ports, tend to face lower trade costs.

Figure 3(a) shows substantial dispersion in trade costs across prefectures, where higher accessibility corresponds to lower costs. Coastal prefectures generally have lower trade costs, shown in darker shading. Even absent the contemporaneous tariff changes following China's WTO accession, exporters in some prefectures face systematically lower trade costs than exporters in others. Section 3.2 and Section 3.3 describe how these three sources of variation are combined into a single measure of prefecture-level trade exposure.

3.2 Construction of the shock

3.2.1 Spatially determined geographic trade costs

Because direct data on prefecture-level, geographically determined trade costs are not available, I construct them in two steps. I first estimate trade costs at the province level in a gravity framework, then downscale them to the prefecture level using each prefecture's proximity to a major seaport.

First, I link the exports from province p in China to the rest of the world (ROW) using a standard gravity equation:

$$x_p^{ROW} = G Y^p Y^{ROW} \phi_p^{ROW} \quad (1)$$

where x_p^{ROW} represents the exports from province p to ROW, and Y^p and Y^{ROW} represent the economic size of the supplier and destination, respectively. Both exports and GDP are measured in 2000. ϕ_p^{ROW} is the bilateral trade accessibility from province p to ROW with $0 \leq \phi_p^{ROW} \leq 1$. To eliminate the constant G and isolate the province-specific accessibility, I normalize the extracted ϕ_p^{ROW} by the national average:

$$\phi_p = \frac{\phi_p^{ROW}}{\frac{1}{P} \sum_{p \in P} \phi_p^{ROW}} \quad (2)$$

where ϕ_p is a scale-free, relative index of trade accessibility. Higher values of ϕ_p indicate

greater accessibility and therefore lower trade costs, reflecting long-lived geographic and infrastructural factors. Throughout the paper, I discuss the underlying friction in terms of trade costs and reserve accessibility for the index itself.

Second, to uncover meaningful within-province heterogeneity, I translate the provincial measure ϕ_p to the prefecture level. For each prefecture i in province p , I compute the minimum great-circle distance to the nearest major seaport, defining $D_i = \ln(\text{distance}_i)$.¹⁶ I then construct an activity-weighted mean distance for the province:

$$\bar{D}_p = \sum_{i \in p} s_{ip} D_i,$$

where s_{ip} is the share of manufacturing employment of prefecture i within province p in 1990. Estimating the relationship between distance and accessibility across provinces yields $\ln \phi_p = \alpha + \lambda \bar{D}_p + u_p$, with a highly significant penalty estimate of $\hat{\lambda} = -0.55$ ($p < 0.01$, $R^2 = 0.54$). Assuming this spatial friction penalty applies homogeneously within provinces, I map the provincial index to prefectures using:¹⁷

$$\ln \phi_i = \ln \phi_p + \hat{\lambda} (D_i - \bar{D}_p), \quad \phi_i = \exp(\ln \phi_i). \quad (3)$$

The term $D_i - \bar{D}_p$ recenters each prefecture's log distance around its province's activity-weighted mean. By construction, the within-province geometric mean of prefectural indices (using the same weights s_{ip}) exactly equals the provincial index, $\prod_{i \in p} \phi_i^{s_{ip}} = \phi_p$. Thus the provincial scale from the gravity step is preserved, and prefectural indices remain comparable across provinces: each ϕ_i equals ϕ_p times a prefecture-specific relative factor $\exp\{\hat{\lambda}(D_i - \bar{D}_p)\}$ whose geometric mean is one.

The resulting measure, ϕ_i , captures prefecture-level, spatially determined geographic trade accessibility before WTO accession, as mapped in Figure 3. In both panels, darker shading indicates lower trade costs, hence higher accessibility and, all else equal, greater exposure to trade shocks. Two patterns emerge. First, Figure 3(a) shows that coastal pre-

¹⁶The exclusive focus on seaports—rather than inland river ports or land border crossings—reflects the logistical realities of China's export boom. According to Chinese Customs Data in 2000, maritime shipping accounted for 63% of the national export value, primarily carrying weight-intensive manufactured goods to the US, the EU, and Japan. Thus, overland distance to a deep-water seaport represents the most binding spatial trade friction. See more discussion in Appendix B.2.2.

¹⁷The downscaling step assumes that the marginal effect of distance on trade costs (λ) is homogeneous across provinces. The penalty for being farther from a port is taken to be the same in inland Gansu as in coastal Jiangsu. This restriction is necessary for mapping the provincial index to prefectures, and the strong fit of the province-level regression supports its plausibility.

fectures generally face lower trade costs than nearby inland prefectures, although some inland areas have relatively good access; for instance, prefectures in Shaanxi have lower trade costs than those in neighboring provinces. Second, Figure 3(b) documents within-province differences in trade costs: prefectures within the same province can differ markedly in market access because of their proximity to major seaports and transport corridors.¹⁸




Figure 3: Geographic Trade Accessibility: Levels and Within-Province Variation.

Notes: Panel 3(a) maps the prefecture-level trade accessibility ϕ_i , with shading based on the national distribution. Panel 3(b) illustrates within-province variation for four provinces, namely Anhui, Zhejiang, Shandong, and Fujian, with shading based on within-province quintiles of trade accessibility and each prefecture's value reported in the subplot legend. In both panels, darker shading indicates lower trade costs and hence higher accessibility. Gray areas denote regions with no data.

3.2.2 Industry-specific trade exposure

To measure the policy-driven component of prefecture-level trade exposure, I weight industry-level tariffs by each prefecture's pre-accession industrial composition, and define it as industry-specific trade exposure of prefecture i :

$$ICS_i = \sum_k \frac{L_{ik,1990}}{L_{i,1990}} \tau_k, \quad (4)$$

¹⁸A visual inspection shows surprisingly high accessibility for certain remote border regions, such as Tibet, mainly due to its exceptionally small GDP; see Appendix C.7 for further discussion. Columns (7)–(8) of Appendix Table A12 show that the core findings are robust to excluding the five major land-border provinces entirely.

where $L_{ik,1990}$ denotes employment in industry k in prefecture i in the base year 1990, $L_{i,1990} = \sum_k L_{ik,1990}$ denotes total employment in prefecture i , and τ_k denotes the tariff level in 2000 of industry k , defined at China's 2-digit CSIC level. A higher value of ICS_i indicates that a larger share of the prefecture's workforce is employed in industries facing larger trade shocks. To construct τ_k , I first compute, for each 4-digit ISIC Rev.3 industry, the average tariff faced by Chinese exporters across China's trading partners in 2000, taken over all partners reporting a tariff for that industry. The tariffs are measured in 2000, before the accession, so they do not respond to the accession itself.¹⁹ I then aggregate these tariffs to the 2-digit CSIC level using a concordance between the two classifications, weighting each ISIC industry by its share of 1990 gross exports within the CSIC industry.

Following the same normalization used for trade accessibility, I isolate the prefecture-specific component by dividing by the national average:

$$\theta_i = \frac{ICS_i}{\frac{1}{I} \sum_{i \in I} ICS_i}, \quad (5)$$

where I is the number of prefectures. A higher θ_i indicates that prefecture i is more exposed to the tariff changes than the national average.

Figure 4(a) maps θ_i , where darker shading denotes greater exposure. Exposure varies markedly across space: coastal prefectures are generally more exposed than inland areas. The dispersion in industry-specific trade exposure is also smaller than that in trade costs.

3.3 Prefecture-specific trade exposure

To construct the composite prefecture-specific trade exposure γ_i , I multiply the trade accessibility index ϕ_i , an inverse measure of trade costs, by the industry-specific trade exposure θ_i .

$$\gamma_i = \phi_i \times \theta_i \quad (6)$$

This multiplicative form is grounded in the Ricardian trade framework (Dornbusch, Fischer and Samuelson, 1977). The core intuition is that the effective price of an exported good in international markets depends jointly on iceberg transportation costs and ad-valorem

¹⁹The object of interest is the tariff reduction industry k experienced between 2001 and 2014. I use the year-2000 tariff level as a proxy for it. Using the realized reduction directly would raise endogeneity concerns, since reductions negotiated multilaterally may reflect political-economic influences and the development of domestic industries. The year-2000 level predates the accession and is strongly correlated with the realized reduction: Appendix Figure A5 shows a correlation of 0.8 across the CSIC manufacturing industries. It therefore captures the prospective tariff reduction each industry would face over the decade following accession, while remaining plausibly exogenous to subsequent developments.




Figure 4: Industry and Prefecture-specific Trade Exposure

Notes: Figure 4(a) shows the industry-specific trade exposure calculated from Equation 5, and Figure 4(b) shows the final prefecture-specific trade exposure calculated from Equation 6. In both panels, darker shading indicates higher trade exposure. Gray areas denote regions with no data.

tariffs. Transportation costs first determine the border price, and ad-valorem tariffs are then levied as a percentage of that price. Because the border price already embeds the spatially determined geographic trade costs, the two barriers compound rather than add, so the total trade barrier is the product of the two. Empirically, this implies that a given level of industry-specific trade exposure generates a larger labour demand shock in a coastal prefecture with low trade costs (a high ϕ_i) than in an isolated inland prefecture facing high trade costs.

Industry-specific trade exposure and trade costs are correlated: export-oriented prefectures tend to have both higher manufacturing concentrations and lower trade costs. They nonetheless provide distinct sources of identifying variation. Prefectures with identical trade costs can face very different trade exposure depending on their predetermined industry mix. Appendix Figure A4 illustrates this at the provincial level, since a prefecture-level plot would be too dense to read. Even among provinces with similar trade costs, industrial composition diverges sharply: Inner Mongolia was historically concentrated in food and textiles, whereas Heilongjiang specialized in wood processing and machinery. The two components are therefore not redundant.

Figure 4(b) maps the distribution of the composite prefecture-specific trade exposure (γ_i). Southeastern coastal prefectures experience the highest overall exposure. Crucially, however, exposure does not decrease monotonically with distance from the coast. Several inland prefectures in provinces such as Ningxia and Shanxi exhibit localized spikes in

exposure, driven by specializations in industries subject to policy changes after the shock.

3.4 Baseline specification

To estimate the effect of trade liberalization on high school enrolment, I employ a cohort difference-in-differences (DiD) design. The baseline specification is estimated at the individual level:

$$y_{jit} = \beta_0 + \beta_1(\gamma_i \times \text{WTO}_t) + \pi_i + \alpha_{pt} + \epsilon_{jit} \quad (7)$$

where y_{jit} is a binary indicator equal to 1 if individual j , residing in prefecture i and turning 16 in year t , ever attended high school, and 0 otherwise. γ_i denotes the continuous, prefecture-specific measure of trade exposure constructed in Section 3.3. The temporal treatment indicator, WTO_t , equals 1 for cohorts turning 16 in 2002 or later, following China’s WTO accession in December 2001, and 0 otherwise.

The coefficient of interest, β_1 , captures how the change in the probability of high school enrolment after WTO accession varies with a prefecture’s exposure to the shock.²⁰ Prefecture fixed effects (π_i) absorb time-invariant local characteristics, such as historical cultural attitudes toward education and baseline geographic endowments. I also include Province $\times$ Year fixed effects (α_{pt}) to absorb time-varying province-specific shocks, such as provincial educational policy rollouts, fiscal spending, and regional macroeconomic trends. All regressions are population-weighted using individual-level census sampling weights. To further account for unobserved shocks that are correlated within local labour markets over time, standard errors are clustered at the prefecture level.

The key identifying assumption is that, in the absence of the trade shock, high school enrolment in prefectures with higher trade exposure would have followed the same trend as in prefectures with lower trade exposure. I assess this assumption by examining pre-trends in the event-study analyses below and conducting additional robustness checks, including an analysis of treatment-effect heterogeneity, in Section 4.6.

To trace the dynamic effects of trade exposure over time, I estimate an event-study ver-

²⁰Trade exposure is a continuous treatment and is strictly positive for all prefectures. Under the strong parallel trends assumption discussed in Section 4.6.3, β_1 recovers a weighted average of marginal responses across exposure levels, where the weights are those implicitly assigned by the two-way fixed effects regression (Callaway, Goodman-Bacon and Sant’Anna, 2026). In that section, I show that these implicit weights do not materially distort the estimand.

sion of the baseline model:

$$y_{jit} = \beta_0 + \sum_{\rho=1995}^{2014} \beta_{\rho}(\gamma_i \times \mathbf{1}\{t = \rho\}) + \pi_i + \alpha_{pt} + \epsilon_{jit} \quad (8)$$

where γ_i is interacted with a full set of year dummies, $\mathbf{1}\{t = \rho\}$, indicating the year in which the cohort turned 16. I omit the interaction term for 2001, the last cohort year before WTO accession. All other fixed effects and the clustering are the same as in Equation 7.

The sequence of estimated coefficients, $\{\beta_{\rho}\}$, serves two purposes. First, for cohorts with $\rho < 2001$, these coefficients test whether enrolment followed similar trends in prefectures with higher and lower trade exposure before WTO accession. Credible identification requires these pre-treatment coefficients to be statistically indistinguishable from zero. Second, for cohorts turning 16 after accession ($\rho > 2001$), these coefficients show how the change in enrolment in year ρ , relative to the 2001 baseline, varies with exposure. Together, these coefficients map the temporal trajectory of the response and allow me to distinguish the immediate response of the earliest affected cohorts from the persistence of the effect across successive cohorts.

4 Results

4.1 Main analysis

I investigate the impact of the trade shock on three outcomes that capture different types of school enrolment. Total high school enrolment captures the extensive margin, and I then separate this enrolment into two tracks: vocational high schools, which provide terminal, job-oriented training, and academic high schools, which prepare students for university. Graduating from an academic high school is in practice the standard route to the National College Entrance Examination (*gaokao*). I therefore define the academic-track outcome inclusively, covering both those whose terminal degree is an academic high school diploma and those who advanced to tertiary education.²¹

Table 1 presents the estimated effects of the WTO trade shock along these three margins: total high school enrolment in Panel A, vocational-track enrolment in Panel B, and academic-track enrolment in Panel C. Column (1) of this table reports the estimates with

²¹ As noted in footnote 10, vocational high schools during this period were generally terminal. In my analysis, any individual with tertiary attainment is classified in the academic track, implying that my estimated effect on vocational high school enrolment represents a conservative lower bound.

Prefecture and Year fixed effects, and Column (2), my baseline specification, adds Province $\times$ Year fixed effects, which absorb unobserved regional macro-trends and local policy roll-outs that vary at the province level over time. In Panel A, it shows that a one-standard-deviation increase in trade exposure reduces the probability of attending high school by 3.46 percentage points ($9.61 \times (-0.0036)$). Given that the baseline high school enrolment rate for the cohort that turned 16 in 2001 was 44%, a one-standard-deviation increase in exposure reduces the likelihood of enrolment by 8% in relative terms ($3.46/44 \approx 8\%$). Panels B and C decompose this aggregate decline by educational track. The estimated coefficients decompose additively up to rounding: $\hat{\beta}_{\text{total}} = \hat{\beta}_{\text{academic}} + \hat{\beta}_{\text{vocational}}$. This holds because academic and vocational enrolment are mutually exclusive and jointly constitute total high school enrolment, and because the specification is linear in the dependent variable with both track regressions using the same sample, weights, and regressors. The decline in total enrolment is driven almost entirely by a sharp, statistically significant decrease in the academic track, while the effect on the vocational track is marginal and statistically insignificant.

The remaining columns of Table 1 progressively add further control variables, and the estimated effects on total and academic enrolment stay stable throughout. Column (3) introduces an “Initial Education $\times$ Year” control, specifically the 1990 middle school completion rate among individuals aged 16 to 18 interacted with year dummies. This allows prefectures with different historical educational endowments to follow distinct trends, and it helps absorb differential educational trajectories stemming from the 985 Project, a national initiative that concentrated university investment in a set of elite institutions and thereby disproportionately favoured regions with historically stronger educational foundations.

Columns (4)–(7) of Table 1 address specific confounding factors, particularly divergent regional trajectories and responses to post-2008 macroeconomic shocks.²² A primary identification concern is that the geographic accessibility measure (ϕ_i) and initial industrial structures might correlate with baseline local wealth and early global integration, which could independently drive distinct educational trends over time. To partial out these unobserved wealth and policy dynamics, Column (4) of Table 1 controls for the interactions between year dummies and two 1995 local macroeconomic variables: log prefecture GDP and the log of Actual Use of Foreign Investment. Because prefecture-level export data were unavailable in 1995, Actual Use of Foreign Investment serves as a proxy for pre-WTO trade

²²In these extended specifications, I replace Province $\times$ Year fixed effects (α_{pt}) with standard Year and Prefecture fixed effects, as Province $\times$ Year fixed effects would absorb much of the identifying variation in these controls needed to identify the coefficient of interest. Appendix Table A13 further shows that results are robust to using Region $\times$ Year fixed effects.

openness (shown in Appendix Figure A6). By absorbing these structural differences, these interactive controls alleviate the concern that the estimated trade effect is confounded by preexisting economic backgrounds or subsequent macro-trends.

Next, following the diagnostics of the shift-share design recommended by [Borusyak, Hull and Jaravel \(2025\)](#), Column (5) of Table 1 controls for the components of the “incomplete share” (ϕ_i and $\sum_k s_{ik}$) interacted with year dummies. Section 4.6.1 discusses this diagnostic in more detail. Finally, in the same table, Columns (6) and (7) account for localized exposure to late-2000s macroeconomic turbulence, particularly the large 2008 state-led fiscal stimulus. Because this domestic stimulus was predominantly channeled through infrastructure projects and state-owned enterprises (SOEs), I control for the 1998 baseline prefecture employment shares in heavy-industrial and infrastructure SOEs, interacted with year dummies: Column (6) adds the heavy-industrial SOE share and Column (7) adds the infrastructure SOE share.²³ These controls partial out the localized labour demand spikes generated by domestic stimulus spending.

Across all seven columns in Table 1, the effect of trade exposure on total and academic high school enrolment remains negative and statistically significant. The effect on vocational enrolment, however, is close to zero, and that in itself is informative. The vocational coefficient is positive and statistically significant in Columns (1) and (4), but it is not stable: it shrinks and loses statistical significance once the geographic and scale controls are added, and is indistinguishable from zero in Columns (5) through (7). The dominant margin of adjustment is therefore between academic schooling and direct labour-force entry, consistent with non-continuation at the end of compulsory schooling rather than a reshuffling across tracks.

4.2 Dynamic impacts and event-study estimates

Figure 5 plots the sequence of estimated event-study coefficients (β_ρ) from Equation 8 across three types of high school enrolment. The plotted models include Prefecture and Year fixed effects, along with interactions between a set of pre-determined baseline controls and year dummies, as detailed in the figure notes.

This event-study figure yields four main results. First, the pre-2001 coefficients for all three outcomes are statistically indistinguishable from zero and exhibit no discernible trend. This finding supports the parallel trends assumption, showing that educational

²³I use the 1998 Chinese Annual Survey of Industrial Firms (ASIF) to construct the baseline prefecture-level employment shares for heavy industries and state-owned enterprises (SOEs).

Table 1: Effect of Trade Exposure on High School Enrolment: Total and by Track

	Baseline Specifications			Extended Controls			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: High School Total							
Trade Exposure	-0.0038*** (0.0006)	-0.0036*** (0.0008)	-0.0031*** (0.0008)	-0.0028*** (0.0006)	-0.0031** (0.0013)	-0.0028** (0.0013)	-0.0028** (0.0013)
Observations	390,377	390,377	390,377	312,072	312,072	312,072	312,072
Adj R-Sq	0.15	0.16	0.16	0.14	0.14	0.14	0.14
Panel B: Vocational High School							
Trade Exposure	0.0007*** (0.0002)	0.0002 (0.0002)	0.0003 (0.0002)	0.0009*** (0.0002)	0.0001 (0.0005)	0.0001 (0.0005)	0.0000 (0.0005)
Observations	390,377	390,377	390,377	312,072	312,072	312,072	312,072
Adj R-Sq	0.02	0.02	0.02	0.02	0.02	0.02	0.02
Panel C: Academic High School							
Trade Exposure	-0.0045*** (0.0005)	-0.0039*** (0.0007)	-0.0034*** (0.0007)	-0.0036*** (0.0005)	-0.0032*** (0.0011)	-0.0028*** (0.0011)	-0.0028** (0.0011)
Observations	390,377	390,377	390,377	312,072	312,072	312,072	312,072
Adj R-Sq	0.13	0.14	0.14	0.12	0.13	0.13	0.13
Prefecture FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes			Yes	Yes	Yes	Yes
Province $\times$ Year FE		Yes	Yes				
Initial Edu $\times$ Year Dummies			Yes	Yes	Yes	Yes	Yes
Base Econ $\times$ Year Dummies				Yes	Yes	Yes	Yes
Geo & Scale $\times$ Year Dummies					Yes	Yes	Yes
Heavy SOE $\times$ Year Dummies						Yes	
Infra SOE $\times$ Year Dummies							Yes

Notes: Robust standard errors in parentheses are clustered at the prefecture level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The table reports estimates from Equation (7), where the outcome is a binary indicator for ever attending the corresponding type of high school by age 16. The sample size drops in Columns (4)–(7) due to missing 1995 prefecture-level GDP and Foreign Investment data. “Initial Edu” is the share of 16–18 year-olds who had completed middle school in the 1990 Population Census. “Base Econ” includes the log of 1995 prefecture GDP and the log of Actual Use of Foreign Investment in 2015 constant USD. “Geo & Scale” comprises the baseline trade cost multiplier (ϕ_i) and the aggregate manufacturing employment share in 1990. “Heavy SOE” and “Infra SOE” denote the 1998 baseline prefecture employment shares in state-owned enterprises in heavy industry and infrastructure sectors, respectively. All regressions are weighted by individual census sampling weights.

trajectories in prefectures experiencing different levels of trade exposure were evolving similarly prior to WTO accession.

Second, a clear divergence emerges after accession: the coefficients for both total and academic high school enrolment begin a decline, becoming statistically significant in 2005. This delayed response reflects both the nature of the outcome variable and the speed of household adjustment. The outcome is a stock measure, capturing all 16-year-olds who ever attended high school, so it responds more slowly than a flow variable would, and households needed time to observe the changing labour market returns before adjusting their schooling choices.

Third, the dynamic estimates reproduce the decomposition established in the baseline

results. The decline in total high school enrolment almost perfectly mirrors the decline in the academic track. The vocational track shows a mild positive movement after accession, but it is statistically insignificant and quantitatively too small to offset the decline in the academic track, consistent with the dominant margin of adjustment being an exit from schooling rather than substitution across tracks.

Finally, the effect on total and academic high school enrolment is not transitory: it persists and widens across successive cohorts over the following decade. China's WTO-related tariff reductions were largely complete by the mid-2000s, so the widening coefficients are unlikely to reflect ongoing reductions in tariffs. Instead, as discussed in Section 4.4, the initial shock was followed by sustained capital investment in labour-intensive export sectors, which is consistent with a more entrenched low-skilled industrial structure in prefectures with higher trade exposure. Each successive cohort thus faced a stronger labour market pull away from academic schooling, not because the shock recurred but because its consequences kept building. This trajectory survives the additional controls, which are intended to absorb the spatially uneven effects of the 2008 fiscal stimulus, the Global Financial Crisis, and concurrent education reforms. Moreover, the core downward trend is already visible in the pre-crisis window (Panel A of Appendix Table A13).

This dynamic parallels the long-lived labour market impacts documented by Autor, Dorn and Hanson (2021) for U.S. local labour markets following the China Shock. However, while persistence in these U.S. labour markets stems largely from the slow depreciation of physical capital and the limited geographic mobility of displaced workers, the persistence I document is consistent with a distinct channel: the difficult-to-reverse schooling decisions of individual students under China's education system. Because the academic and vocational tracks are separated at the point of high school entry, students who exit the academic pipeline in response to short-run labour market incentives face severely constrained pathways back into the academic track, and with it, into university education. This effect accumulates over the sample period and plausibly reshapes the skill composition of successive local cohorts.

4.3 Heterogeneous effects

The main results show a negative effect of the trade shock on high school enrolment for China as a whole. Because local labour market frictions and household constraints vary across China's regions and demographic groups, this single estimate may conceal considerable variation. I therefore examine how the effect differs by gender and region.




Figure 5: Dynamic Effect of Trade Exposure on High School Enrolment: Total and by Track

Notes: The figure plots event study estimates of trade exposure on three high school enrolment outcomes: total high school enrolment, academic high school enrolment, and vocational high school enrolment. The dashed vertical line denotes 2001, the year before China's accession to the WTO. All models include Prefecture and Year FE, along with interactions between year dummies and baseline controls (economic scale, initial middle school completion rates, and heavy SOE employment shares). Dots represent point estimates, and error bars indicate 90% confidence intervals using standard errors clustered at the prefecture level. All regressions are weighted by individual census sampling weights.

Columns (1)–(2) of Table 2 show that the negative effect is concentrated among boys. Boys experience a large reduction in academic enrolment with no offsetting vocational gain, while the point estimate for girls is smaller and statistically insignificant. This aligns with the gender-biased nature of post-WTO labour demand, since the manufacturing boom drew heavily on physical and blue-collar labour, and [Li et al. \(2019\)](#) find that the effect of the China Shock on years of schooling is likewise concentrated among young men.

Columns (3)–(6) of Table 2 show that the effect of the trade shock on high school enrolment varies across regions; the provinces in each region are listed in the table notes. The East is the only region where the negative effect on both total and academic enrolment is statistically significant. This echoes the geography of labour market access documented by [Facchini et al. \(2019\)](#): for students in the East region, export factories are often located in or near their home prefectures, making the opportunity cost of schooling salient. Outside the East, none of the estimates is statistically significant. The point estimates nonetheless show a pattern worth noting. The Central region shows negative but insignificant estimates across all three margins; its inland location likely placed the coastal factory jobs that drive the Eastern effect largely out of reach for adolescents at the point of the enrolment decision. Nationally, census data from 2005 show that inter-provincial migration differed by

age: 11.1% among young adults aged 19–30 but only 2.5% among minors aged 15–17. The Western region shows positive but imprecisely estimated effects on total enrolment, and the point estimates suggest a shift toward the vocational track, consistent with a household investment channel in which remittances from migrating family members ease liquidity constraints and redirect resources toward vocational training (De Brauw and Giles, 2017). The Northeast shows small and insignificant coefficients across all panels. Long dominated by state-owned enterprises and heavy industry, this region was largely bypassed by a shock that primarily stimulated light manufacturing, which may have left the opportunity cost of schooling largely unchanged.

Table 2: Heterogeneous Effects of Trade Exposure on High School Enrolment by Gender and Region

	Gender		Region			
	(1) Boys	(2) Girls	(3) East	(4) Central	(5) West	(6) Northeast
Panel A: High School Total						
Trade Exposure	-0.0040*** (0.0012)	-0.0016 (0.0017)	-0.0029* (0.0015)	-0.0252 (0.0485)	0.0197 (0.0370)	0.0061 (0.0251)
Observations	158,025	154,047	160,574	82,318	45,866	23,314
Adj R-Sq	0.13	0.17	0.13	0.16	0.18	0.14
Panel B: Vocational High School						
Trade Exposure	-0.0002 (0.0005)	0.0004 (0.0007)	0.0001 (0.0005)	-0.0243 (0.0258)	0.0354 (0.0373)	-0.0013 (0.0075)
Observations	158,025	154,047	160,574	82,318	45,866	23,314
Adj R-Sq	0.02	0.02	0.02	0.02	0.02	0.01
Panel C: Academic High School						
Trade Exposure	-0.0037*** (0.0012)	-0.0020 (0.0013)	-0.0030** (0.0014)	-0.0009 (0.0553)	-0.0157 (0.0598)	0.0073 (0.0242)
Observations	158,025	154,047	160,574	82,318	45,866	23,314
Adj R-Sq	0.11	0.15	0.11	0.15	0.16	0.14

Notes: Robust standard errors in parentheses are clustered at the prefecture level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The regressions use the control set from Column (7) of Table 1. The regional classification follows the standard four-region division used in Chinese official statistics: East (Beijing, Tianjin, Hebei, Shanghai, Jiangsu, Zhejiang, Fujian, Shandong, Guangdong, Hainan), Central (Shanxi, Anhui, Jiangxi, Henan, Hubei, Hunan), West (Inner Mongolia, Guangxi, Chongqing, Sichuan, Guizhou, Yunnan, Tibet, Shaanxi, Gansu, Qinghai, Ningxia, Xinjiang), and Northeast (Liaoning, Jilin, Heilongjiang). All regressions are weighted by individual census sampling weights.

4.4 Interpreting the persistence: capital-labour complementarity

A notable feature of the event-study results in Figure 5 is that the negative effect of the trade shock on total and academic high school enrolment persists and widens throughout

the subsequent decade. To investigate this persistence, I draw on the complementarity between fixed capital investment and production workers (Curtis et al., 2021). Using data from 2000 to 2010, a period over which the capital and employment measures are consistently defined, I examine the relationship between trade exposure, capital accumulation (measured by net and total fixed assets), and labour demand.²⁴

I first document the complementarity at the industry level. Appendix Figure A7 shows that both net and total fixed assets are strongly positively correlated with employment ($R^2 > 0.40$). This complementarity is most pronounced in China's leading labour-intensive export sectors, such as textiles, apparel, and electronics assembly. Large-scale capacity expansions in these sectors therefore draw in substantial numbers of manual, low-skilled workers.

I then turn to the spatial evidence to show how this capital-labour pattern maps onto the geographic distribution of the China Shock. Appendix Figure A8 shows the relationship between investment and employment across provinces, with the size and shading of the markers representing the intensity of the trade shock at the province level. The complementarity between production capacity and labour demand also holds at the province level, and provinces with greater trade exposure show disproportionately higher levels of both fixed asset investment and total employment. Rather than a uniform national industrial boom, the expansion in production capacity was geographically concentrated along the highly exposed coastal provinces.

Together, these patterns describe a channel that could account for the persistent decline in high school enrolment. Because the capital accumulated in these labour-intensive sectors is complementary to low-skilled labour, the demand it sustained did not fade once the shock had passed. This industrial capacity is consistent with a continued pull on adolescents away from the academic track, so that a one-time trade policy change had sustained consequences for high school enrolment across cohorts within my sample window.

4.5 Distributional impacts: cohort-specific educational inequality

Having shown that the negative effect of the trade shock on total high school enrolment is persistent, I next examine whether this long-lasting effect reshaped the distribution of educational attainment within affected cohorts, and whether any such reshaping changed educational inequality. Measuring this inequality correctly requires isolating the cohorts

²⁴The capital data are available only at the province level, so this evidence is coarser than in the baseline analysis. I therefore treat it as consistent with the proposed channel rather than as a direct test of it.

actually exposed to the shock. If one simply computed inequality over all residents of a prefecture, the measure would be contaminated by older people who made their decisions long before accession. To avoid this, my measure of educational inequality is cohort-by-cohort, capturing only the final attainment of individuals born in a given year.

I restrict the sample to birth cohorts 1979–1995: the upper bound ensures that the youngest individuals in the sample were old enough by 2014 to have entered any tertiary education they would pursue, so that their final attainment is not understated. I sort individuals into six ordered attainment categories, from no schooling to graduate level, and compute an education Gini index ($G_{it}^{\text{edu}} \in [0, 1]$) for each prefecture-cohort cell using census sampling weights. The full calculation of the education Gini index is in Appendix B.4. To ensure these estimates are not driven by confounders, I use the control set from Column (7) of Table 1 to absorb concurrent educational reforms. Additionally, Appendix Table A4 reconstructs the measures on a non-migrant sample based on *hukou* registration to rule out selective migration.

In Columns (1)–(2) of Table 3, I present the impact of the trade shock on the education Gini index. With fixed effects only, Column (1) shows a negative and significant effect, meaning that attainment appears more equally distributed. Once the controls are added in Column (2), the estimate becomes small and statistically insignificant. The negative estimate in Column (1) therefore appears to reflect pre-existing regional characteristics correlated with trade exposure, such as local economic scale, rather than the trade shock itself. Column (2) provides no evidence of a net compression or widening of local educational inequality.

To understand the compositional changes behind this stability in overall inequality, I group individuals into three categories by highest attainment, namely middle school or below, high school, and tertiary or above, and measure each cohort’s share in each category. The results, in Columns (3)–(8) of Table 3, point to a restructuring of the distribution rather than a uniform shift. Trade exposure significantly raises the share of each cohort whose highest attainment is middle school, the final stage of China’s compulsory education, and significantly lowers the share that attains tertiary education. The impact of the shock on the high school share is less robust: its coefficient changes sign and loses statistical significance once the controls are added, so it should be read with caution. Because this expansion at the bottom (stopping at middle school) and contraction at the top (forgoing university) occur at opposite ends of the distribution, they could push measured inequality in opposite directions, which is consistent with the result in Column (2) in the same table.

This restructuring truncates the upper end of the academic pipeline at two possible

points. Some students may not enter academic high school at all; others may finish academic high school but forgo university, given the cost of college and the high opportunity cost of continued schooling in exposed regions. The unstable coefficient on the high school terminal share is consistent with both margins being at work. This narrowing of tertiary attainment may make skill upgrading harder for high-exposure regions if they seek to move away from labour-intensive manufacturing.

Table 3: Trade Exposure, Educational Inequality, and Attainment Distribution

	Gini Index		Middle School		High School		Tertiary	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Trade Exposure	-0.0003*** (0.0001)	0.0003 (0.0002)	0.0025*** (0.0007)	0.0025** (0.0011)	-0.0004* (0.0002)	0.0011 (0.0007)	-0.0024*** (0.0007)	-0.0038*** (0.0011)
Observations	5,593	3,859	5,593	3,859	5,593	3,859	5,593	3,859
Adj R-Sq	0.67	0.50	0.76	0.81	0.47	0.47	0.80	0.83
Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Baseline Controls		Yes		Yes		Yes		Yes

Notes: Robust standard errors in parentheses are clustered at the prefecture level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Odd columns include prefecture and province $\times$ year fixed effects only, and even columns instead use prefecture and year fixed effects with controls defined in Column (7) of Table 1; see the notes to that table for detailed definitions of these controls. The four outcome variables, the Gini Index and the shares whose highest attainment is middle school, high school, or tertiary education, are calculated at the prefecture-cohort level using census sampling weights, and all regressions are weighted by each cell's census-weighted population.

4.6 Robustness checks

I conduct four robustness checks: validating the modified shift-share design, testing sensitivity to violations of parallel trends, exploring treatment-effect heterogeneity, and addressing internal migration. Additional checks, including placebo tests and alternative measures of trade exposure, are in Appendices C and D.

4.6.1 Validity of the modified shift-share design

Justification via the shifts-based framework. A concern with incorporating the trade cost multiplier (ϕ_i) into the classic shift-share measure (θ_i) is that it is partially derived from baseline provincial GDP and export data, so the composite exposure measure γ_i is mechanically higher in richer provinces. Since these provinces also invest more in public education, which could raise enrolment, omitting local income pushes the estimated coefficient of the trade shock on high school enrolment upward. Because this coefficient

is negative, the omission biases the baseline estimates toward zero and thus makes them conservative. To establish the validity of this identification strategy, I follow the checklist of [Borusyak, Hull and Jaravel \(2025\)](#).

The overall composite trade exposure decomposes algebraically as

$$\gamma_i = \phi_i \times \theta_i = \phi_i \times \sum_k \left(\frac{1}{M} s_{ik} \right) \tau_k = \sum_k \underbrace{\left(\frac{1}{M} \phi_i s_{ik} \right)}_{\text{Effective Share}} \tau_k,$$

τ_k represents the industry-level tariff (*shifts*), $\frac{1}{M} \phi_i s_{ik}$ is the geography-weighted effective exposure (*share*), and M denotes the national average normalization scalar. Causal identification requires the shifts to be quasi-randomly assigned and orthogonal to local unobservables, and permits the shares to be endogenous provided their structural components are partialled out.

Establishing this involves four requirements. First, the shifts themselves must be exogenous: the baseline 2000 tariffs and their subsequent cuts were set through multilateral negotiations and imposed by China’s trading partners, so they are plausibly exogenous to the educational and labour market decisions of any single prefecture. Second, the framework requires controlling for the “incomplete share”. Without this control, prefectures with larger manufacturing sectors or lower trade costs would mechanically receive greater exposure, and that relationship could itself correlate with local educational trends. This incomplete share in my case is the sum of effective shares ($\frac{1}{M} \phi_i \sum_k s_{ik}$), the product of the trade cost multiplier (ϕ_i) and baseline manufacturing scale ($\sum_k s_{ik}$). I control for each component separately, interacted with year dummies, which spans this composite term under a less restrictive functional form, so that the remaining identifying variation rests purely on the interaction of local industrial composition with the exogenous tariff cuts. Third, in terms of the effective sample size of the shifts and their exogeneity, I show that across the 38 manufacturing industries in my sample, the effective number of shifts is 16.46, calculated as the inverse of the Herfindahl index, indicating that identification draws on a reasonably broad distribution of quasi-random shocks rather than being driven by a single sector. Appendix C.1.1 presents a formal balance test showing that industry-level tariffs in 2000 are statistically uncorrelated with 1990 pre-determined industry characteristics. Fourth, I implement exposure-robust standard errors following [Borusyak, Hull and Jaravel \(2025\)](#) to account for correlated trade shocks across regions with similar industrial compositions, and the negative effect of the trade shock on high school enrolment remains statistically significant. The industry-level equivalent regression and the finite-sample dis-

cussion are reported in Appendix C.1.2.

Absorbing macroeconomic confounders. Columns (4)–(7) of Table 1 progressively add the macroeconomic controls. Because WTO accession coincided with improvements in contract enforcement and property rights protection in China, Panel B of Appendix Table A13 further controls for baseline contract intensity interacted with year dummies, following [Facchini et al. \(2019\)](#) and [Erten and Leight \(2021\)](#).²⁵ Across these specifications, the trade exposure coefficient remains similar in magnitude and significance.

Alternative measurement using pure physical geography. To further alleviate concerns that ϕ_i introduces endogenous economic factors into the treatment, I replace the gravity-derived trade shock with the inverse of each prefecture’s physical distance to the nearest major deep-water seaport operational as of 1990, normalized by the national average. This distance is time-invariant, and because the port network predates the accession, as described in Section 2.2, it cannot have been shaped by the subsequent export surge or by policy favouritism toward particular prefectures. As reported in Appendix D.1, this measure yields results similar to the baseline in sign and magnitude across all specifications.

4.6.2 Sensitivity to the parallel trends assumption

The cohort difference-in-differences design assumes that, absent WTO accession, enrolment in prefectures with different levels of exposure would have evolved in parallel. I assess this assumption in three ways.

Event-study evidence. As shown in Figure 5, the pre-accession coefficients are statistically indistinguishable from zero, consistent with the assumption that prefectures with different levels of exposure were on parallel trajectories prior to WTO entry.

Bounding exercise. I implement the relative magnitudes restrictions of [Rambachan and Roth \(2023\)](#), which quantify how large a post-treatment violation would need to be, relative to the largest pre-treatment deviation, before the confidence interval no longer excludes zero. Figure 6 presents the 90% confidence intervals for the average treatment effect over the post-crisis period (2009–2014). For both total and academic high school enrolment, the results are robust to violations up to $\bar{M} \approx 0.5$, that is, to post-treatment violations up to half the largest pre-treatment deviation. The vocational-track estimate is not significant under exact parallel trends, so its breakdown value is zero. Details on this method and results with additional controls are in Appendix C.2.

²⁵The contract intensity index follows [Nunn \(2007\)](#); construction details are in Appendix C.6.




Figure 6: Sensitivity Analysis: Later-Period Effect (2009–2014)

Notes: The figure plots 90% robust confidence intervals under the relative magnitudes restriction $\Delta^{RM}(\bar{M})$ of [Rambachan and Roth \(2023\)](#), which bounds the post-treatment violation of parallel trends by $\bar{M}$ times the largest pre-treatment deviation; $\bar{M} = 0$ corresponds to exact parallel trends. The target parameter is the effect over the later post-treatment window (2009–2014), estimated from Equation 8 with prefecture and year fixed effects. All regressions are weighted by individual census sampling weights.

Isolating contemporaneous macroeconomic shocks. Beyond the macroeconomic controls already included in the main results, I further isolate the early post-WTO effects from late-2000s macroeconomic turbulence through two complementary sample restrictions. First, Columns (5)–(6) of Appendix Table A12 exclude the period 2009 onward, showing that the decline in academic enrolment is firmly underway before the onset of macroeconomic volatility. Second, Panel A of Appendix Table A13 splits the post-treatment window into an early (2002–2007) and a later (2008–2014) period: the negative effect is still statistically significant in the pre-crisis window 2002–2007, showing that the core result is not driven by post-2008 disruptions.

4.6.3 Continuous exposure and treatment-effect heterogeneity

In designs with a continuous treatment, the standard parallel-trends assumption is not sufficient to give the coefficient of interest β_1 a causal interpretation as an average marginal response. As [Callaway, Goodman-Bacon and Sant’Anna \(2026\)](#) show, identification requires *strong parallel trends*, which requires that the change in enrolment observed at a given exposure level reflect what would have occurred for prefectures at every exposure level had they been exposed to that same level. Because these additional requirements involve potential outcomes under different exposure levels, they cannot be tested directly using conventional pre-trend checks. Following their recommendations, I provide four complementary pieces of indirect evidence, reported in detail in Appendix C.3, that the

estimates are not driven by treatment-effect heterogeneity across exposure levels.

First, the strong parallel trends assumption becomes more plausible after conditioning on observed covariates (Callaway, Goodman-Bacon and Sant’Anna, 2026). The specifications in Columns (3)–(7) of Table 1 already include baseline characteristics interacted with year dummies, representing the dimensions along which prefectures both selected into exposure and might have differed in their response.

Second, I redefine the treatment as a binary indicator equal to one for the top quintile of trade exposure and zero otherwise. I interact this binary indicator with year dummies and estimate an event-study specification including the same set of controls as I use for Figure 5. Appendix Figure A10 shows that this definition yields a similar pattern of negative effects on total and academic high school enrolment across cohorts.

Third, I verify that the estimated coefficient is not distorted by the weighting that a continuous-treatment two-way fixed-effects regression implicitly assigns across exposure levels (Callaway, Goodman-Bacon and Sant’Anna, 2024). I compute the global average causal response ($ACRT^{glob}$) of Callaway, Goodman-Bacon and Sant’Anna (2026), which averages the marginal effect of a further increase in exposure, defined at each exposure level across all exposed prefectures, over the empirical distribution of exposure. As shown in Appendix Table A7, the results are nearly identical to those in Column (6) of Table 1.

Fourth, I examine directly whether the marginal response to exposure is constant, since heterogeneity would show up as a nonlinear exposure-response relationship. I divide the sample into 15 equally sized groups by trade exposure, a partition fine enough to trace the shape of the relationship across the full support of exposure. Appendix Figure A11 plots the change in enrolment between the pre- and post-accession periods (y-axis) against the level of trade exposure (x-axis) for each group. The relationship shows no sign of curvature, suggesting that the marginal response does not vary sharply with the level of exposure. The pattern is preserved when the five most exposed prefectures are excluded, as shown in Appendix Figure A12 and Appendix Table A8.

4.6.4 Addressing internal migration and spatial sorting

One concern in my analysis is whether the observed educational declines reflect changes in schooling decisions or merely demographic compositional changes driven by internal migration. If prefectures with higher trade exposure attract an influx of low-educated migrant workers, or if local trade shocks induce out-migration of youth, the observed declines could instead reflect a shift in local demographics. I address this concern in three steps.

Demographic baseline. In Appendix C.5.2 I show that youth mobility across prefectures was constrained during the high-school decision window (ages 16–18): between 85% and 93% of youth remained in their home prefecture across the three five-year windows spanning 2000 to 2015. Intra-provincial youth migrants consistently exhibited higher average schooling than local non-migrants, consistent with intra-provincial mobility being driven by relocation toward prefecture capitals for better school access rather than toward manufacturing employment.

Isolating pure locals and controlling for spatial spillovers. Appendix Table A11 presents three complementary tests using the 2015 Census to show that migration does not drive the estimated decline. First, a placebo-style check using youth out-migration as the outcome variable finds no significant effect of local trade exposure on the probability of out-migration. Second, following Tombe and Zhu (2019), I compare each individual’s current prefecture of residence with their registered *hukou* prefecture recorded in the Census to construct a pure locals sample by excluding all inter- and intra-provincial cross-prefecture migrants. The results remain negative and statistically significant, with similar magnitudes. Third, following the measurement in Li (2018), adding spatial spillover controls for neighbouring prefectures and distance-weighted trade exposure of all other prefectures still yields a negative effect of the trade shock on high school enrolment, while both spillover terms are statistically insignificant, consistent with the estimates capturing localized behavioural responses rather than the pull of distant labour market opportunities. Construction details and full results are reported in Appendix C.5.3.

Addressing selective early migration using longitudinal data. The Census-based pure-locals restriction mitigates bias from recent and adult migration but cannot resolve spatial sorting that occurred prior to age 16. To address this, I turn to the 2018 wave of China Family Panel Studies (CFPS), a nationally representative longitudinal survey that asks respondents which province they resided in at birth and at age 12—a retrospective anchor that predates any independent, labour-driven migration decision. Under China’s strict *hukou* system and school enrolment rules as I discussed in Section 2.1, a student’s age-12 location effectively determined the province in which they sit the high school entrance exam (*zhongkao*) and enrol, so matching individuals to this location maps them to trade shocks that are plausibly exogenous to their subsequent educational choices. In Table 4, I assign exposure by province at birth and by province at age 12; the results are consistent with the core finding under both definitions. See Appendix C.5.4 for a description of the data sources and further discussion.

Table 4: Robustness: Addressing Selective Early Migration (CFPS Data)

	Province at Age 12		Province at Birth	
	HS Total		HS Total	
	(1)	(2)	(3)	(4)
Trade Exposure	-0.0068*** (0.0021)	-0.0122*** (0.0032)	-0.0052*** (0.0019)	-0.0103*** (0.0031)
Observations	6,943	6,943	6,867	6,867
Adj R-Sq	0.276	0.277	0.276	0.277
Baseline controls	Yes	Yes	Yes	Yes
Geo & Scale $\times$ Year Dummies		Yes		Yes

Notes: Robust standard errors in parentheses are clustered at the province level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Trade exposure is assigned at the province level, and the dependent variable is a binary indicator for ever attending any type of high school by age 16. “Baseline controls” include province fixed effects, region $\times$ year fixed effects, gender, *hukou* status at age 12 (agricultural versus non-agricultural *hukou*), and the 1990 baseline middle school completion rate interacted with year dummies. Columns (2) and (4) also control for trade accessibility and the manufacturing employment share, both interacted with year dummies. All regressions are weighted by the CFPS 2018 individual cross-sectional sampling weights.

5 Conclusion

This paper examines the impact of China’s accession to the WTO in December 2001 on high school enrolment, using a new measure of prefecture-level trade exposure that incorporates spatially determined geographic trade costs in addition to industrial composition and tariff reductions. I find that a one-standard-deviation increase in trade exposure reduces high school enrolment by 8% relative to the baseline enrolment rate of the last pre-accession cohort, the cohort that turned 16 in 2001. This decline is driven almost entirely by a sharp contraction in academic high school enrolment, and it persists and grows across successive cohorts over the subsequent decade. This persistence is consistent with sustained capital investment in labour-intensive export sectors, which would have kept the opportunity cost of schooling elevated. Overall educational inequality is largely unchanged, but the educational distribution in high-exposure regions is reshaped from within: fewer students reach tertiary education and more stop at the compulsory level. This is a recomposition of attainment rather than a broad-based upgrading. In this paper, I use enrolment as a proxy for education, so these results speak to the quantity of schooling rather than its quality, which is difficult to measure consistently across prefectures over this period.

In high-exposure regions, export-oriented growth after the accession plausibly raised the demand for low-skilled labour, discouraging continued schooling. This paper suggests a tension between short-run labour market gains and the long-run educational attainment of successive cohorts. Whether this pattern generalizes beyond China remains an open

question, but the channel suggested here, sustained capital investment in labour-intensive sectors, could be relevant for other developing economies integrating into global markets. As automation increasingly substitutes for routine labour, cohorts drawn out of school in this way may be concentrated in the jobs most exposed to automation, raising the long-run cost of the schooling response.

References

- Acemoglu, Daron, and Pascual Restrepo.** 2022. "Tasks, automation, and the rise in US wage inequality." *Econometrica*, 90(5): 1973–2016.
- Anderson, James E, and Eric Van Wincoop.** 2004. "Trade costs." *Journal of Economic literature*, 42(3): 691–751.
- Atkin, David.** 2016. "Endogenous skill acquisition and export manufacturing in Mexico." *American Economic Review*, 106(8): 2046–2085.
- Autor, David, David Dorn, and Gordon H Hanson.** 2021. "On the persistence of the China shock." *National Bureau of Economic Research*.
- Blanchard, Emily J, and William W Olney.** 2017. "Globalization and human capital investment: Export composition drives educational attainment." *Journal of International Economics*, 106: 165–183.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel.** 2025. "A practical guide to shift-share instruments." *Journal of Economic Perspectives*, 39(1): 181–204.
- Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna.** 2024. "Event Studies with a Continuous Treatment." *AEA Papers and Proceedings*, 114: 601–605.
- Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna.** 2026. "Difference-in-Differences with a Continuous Treatment." *American Economic Review*. Forthcoming.
- Curtis, E Mark, Daniel G Garrett, Eric C Ohrn, Kevin A Roberts, and Juan Carlos Suárez Serrato.** 2021. "Capital investment and labor demand." National Bureau of Economic Research.
- De Brauw, Alan, and John Giles.** 2017. "Migrant opportunity and the educational attainment of youth in rural China." *Journal of Human Resources*, 52(1): 272–311.
- Ding, Weili, and Steven F Lehrer.** 2007. "Do peers affect student achievement in China's secondary schools?" *The review of economics and statistics*, 89(2): 300–312.
- Dornbusch, Rudiger, Stanley Fischer, and Paul Anthony Samuelson.** 1977. "Comparative advantage, trade, and payments in a Ricardian model with a continuum of goods." *The American Economic Review*, 67(5): 823–839.
- Edmonds, Eric V, and Nina Pavcnik.** 2005. "The effect of trade liberalization on child labor." *Journal of international Economics*, 65(2): 401–419.
- Edmonds, Eric V, Nina Pavcnik, and Petia Topalova.** 2010. "Trade adjustment and human capital investments: Evidence from Indian tariff reform." *American Economic Journal: Applied Economics*, 2(4): 42–75.

- Erten, Bilge, and Jessica Leight.** 2021. "Exporting out of agriculture: The impact of WTO accession on structural transformation in China." *Review of Economics and Statistics*, 103(2): 364–380.
- Facchini, Giovanni, Maggie Y Liu, Anna Maria Mayda, and Minghai Zhou.** 2019. "China's "Great Migration": The impact of the reduction in trade policy uncertainty." *Journal of International Economics*, 120: 126–144.
- Hanushek, Eric A, and Ludger Woessmann.** 2008. "The role of cognitive skills in economic development." *Journal of economic literature*, 46(3): 607–668.
- Jiang, Xuan, Kendall Kennedy, and Jiatong Zhong.** 2023. "When Opportunity Knocks: China's Open Door Policy and Declining Educational Attainment." *Labour Economics*, 81: 102312.
- Lehrer, Steven F., and Peter Lin Qi Wang.** 2026. "The Returns to Cognitive and Noncognitive Skills in Early-Career Labor Markets: Evidence from Gansu, China." SSRN Working Paper No. 6876904.
- Li, Bingjing.** 2018. "Export expansion, skill acquisition and industry specialization: Evidence from China." *Journal of International Economics*, 114: 346–361.
- Li, Hongbin, Prashant Loyalka, Scott Rozelle, and Binzhen Wu.** 2017. "Human capital and China's future growth." *Journal of Economic Perspectives*, 31(1): 25–48.
- Li, Jie, Yi Lu, Hong Song, and Huihua Xie.** 2019. "Long-term impact of trade liberalization on human capital formation." *Journal of Comparative Economics*, 47(4): 946–961.
- Liu, Chengfang, Linxiu Zhang, Renfu Luo, Scott Rozelle, Brian Sharbono, and Yaojiang Shi.** 2009. "Development challenges, tuition barriers, and high school education in China." *Asia Pacific Journal of Education*, 29(4): 503–520.
- Li, ZR, RD Yang, and GT Huang.** 2022. "Intra-national trade costs, international trade costs and the promotion of the domestic circulation of China." *China Industrial Economics*, 2: 61–79.
- Ministry of Education of the People's Republic of China.** 2001. "Statistical Report on National Educational Development in 2000." *Ministry of Education of the People's Republic of China*, In Chinese. Available at the Ministry of Education website.
- Nunn, Nathan.** 2007. "Relationship-specificity, incomplete contracts, and the pattern of trade." *The quarterly journal of economics*, 122(2): 569–600.
- Pierce, Justin R, and Peter K Schott.** 2016. "The surprisingly swift decline of US manufacturing employment." *American Economic Review*, 106(7): 1632–1662.
- Rambachan, Ashesh, and Jonathan Roth.** 2023. "A more credible approach to parallel trends." *Review of Economic Studies*, 90(5): 2555–2591.
- Tombe, Trevor, and Xiaodong Zhu.** 2019. "Trade, migration, and productivity: A quantitative analysis of china." *American Economic Review*, 109(5): 1843–1872.
- Zhang, Junsen, and Kang Zhou.** 2023. "Quota removal, destination-specific export shocks, and skill acquisition in China." *Journal of Development Economics*, 165: 103149.
- Zhang, Tiedao, and Minxia Zhao.** 2006. "Universalizing nine-year compulsory education for poverty reduction in rural China." *International review of education*, 52(3): 261–286.