The Importance of Educational Credentials: Schooling Decisions and Returns in Modern China

Alex Eble and Feng Hu*

May 2016

Abstract

All governments who provide primary and secondary education must choose the number of years of schooling needed to earn educational credentials such as a high school diploma. We exploit a policy change in China to estimate the impact of this choice on schooling decisions and labor supply. The policy extended the length of primary school by one year and was rolled out gradually across China over 25 years. We show this has reallocated 850 billion person-hours from labor to schooling to date, as the vast majority of affected individuals chose to earn a credential despite the additional year in school this required. We use this result to estimate the labor market returns to an additional year of schooling, holding highest credential constant. We show that the year generates small average returns in the labor market, but is highly redistributive towards China's disadvantaged. We then provide evidence that our estimate is very close to a parameter of central interest: the labor market returns to the human capital accumulated in a year of school, isolated from the contribution of signaling. A cost-benefit exercise estimates that the policy, while redistributive, generated a likely net loss of tens of billions of dollars. We conclude with evidence that some of results are generalizable to other contexts in the developing world.

*Eble (corresponding author): Teachers College, Columbia University. Email: eble@tc.columbia.edu; Hu: Dongling School of Economics and Management, University of Science and Technology Beijing. Email: feng3hu@gmail.com. We would like to thank Andrew Foster, Emily Oster, and John Tyler for extensive feedback and guidance, and Alexei Abrahams, Anna Aizer, Dionissi Aliprantis, Marianna Battaglia, Natalie Bau, Nate Baum-Snow, Ken Chay, Andrew Elzinga, John Friedman, David Glancy, Nate Hilger, Rob Jensen, Melanie Khamis, Eoin McGuirk, Bryce Millett-Steinberg, Sri Nagavarupu, Gareth Olds, Anja Sautmann, Rajiv Sethi, Jesse Shapiro, Zach Sullivan, Felipe Valencia, and many seminar audiences for helpful suggestions. Eble acknowledges support from the US National Science Foundation through Graduate Research and IGERT Fellowships, and from the Brown University PSTC. Hu acknowledges financial support from the National Natural Science Foundation of China (71373002, 71133003, 71420107023). All remaining errors are our own. JEL codes: I25, I26, J24.

I. Introduction

A crucial role of government is to provide and regulate schooling, and a key decision governments face is to determine the number of years needed to complete primary and secondary education. This choice of "credential length" also generates an important empirical regularity: across the majority of developing and developed countries, most individuals leave school and enter the labor market after completing a given level of school and earning a credential, such as a middle or high school diploma[1]. Credential length policy thus determines how most of the world's people spend the early productive years of their lives. Furthermore, this is an active policy area, with recent changes to credential length in several developing countries.

In this paper, we exploit a policy change in China to estimate this type of policy's impact on schooling decisions and labor supply. In addition, we use this as variation to estimate three key parameters characterizing the demand for educational credentials in modern-day China, the returns to schooling in the Chinese labor market that accrue through human capital accumulation as opposed to signaling, and the public finance implications of the policy we study.

In 1980 the Chinese government announced that it would increase by one the number of years needed to complete primary school while leaving unchanged the national curriculum and length of all other levels of schooling. This policy was rolled out gradually across localities over 25 years and has induced over 400 million people to spend an additional year in primary school so far. The policy was implemented as a part of efforts to prepare China's labor force for the shift from a command to a market economy.

The policy went into effect in each locality at a time when the modal student (also the median) left school after spending eight years in school, five in primary and three in middle school, and receiving her middle school diploma. The policy changed the parameters of the schooling decision individuals faced. Now eight years of schooling meant six years in primary, two in middle, and no middle school credential. To keep credential attainment constant, our modal student would have to spend an additional year in school and thus out of the labor market[2]. Given the bunching at credential attainment years we observe prior to the policy, more than 75% of students in our sample faced a similar decision, either at the middle school, high school, or tertiary margin. Figure I depicts the Chinese education system before and after the policy change and this modal student's post-policy schooling decision.

[1] According to the 2014 Current Population Survey, approximately three quarters of US citizens aged 30-64 finish their schooling in a credential attainment year. In our data from China, this fraction is more than 80%. We show in Section VI that similar ratios present in 48 of the 74 countries in which we have Demographic and Health Surveys data on completed years of schooling.

[2] Though middle school was made compulsory in 1986, in Appendix 3 we show evidence that the rollout of the compulsory middle school policy has little effect on whether or not individuals complete middle school or earn at least a middle school credential.

Figure I: Years of schooling to earn credentials, pre- and post-policy

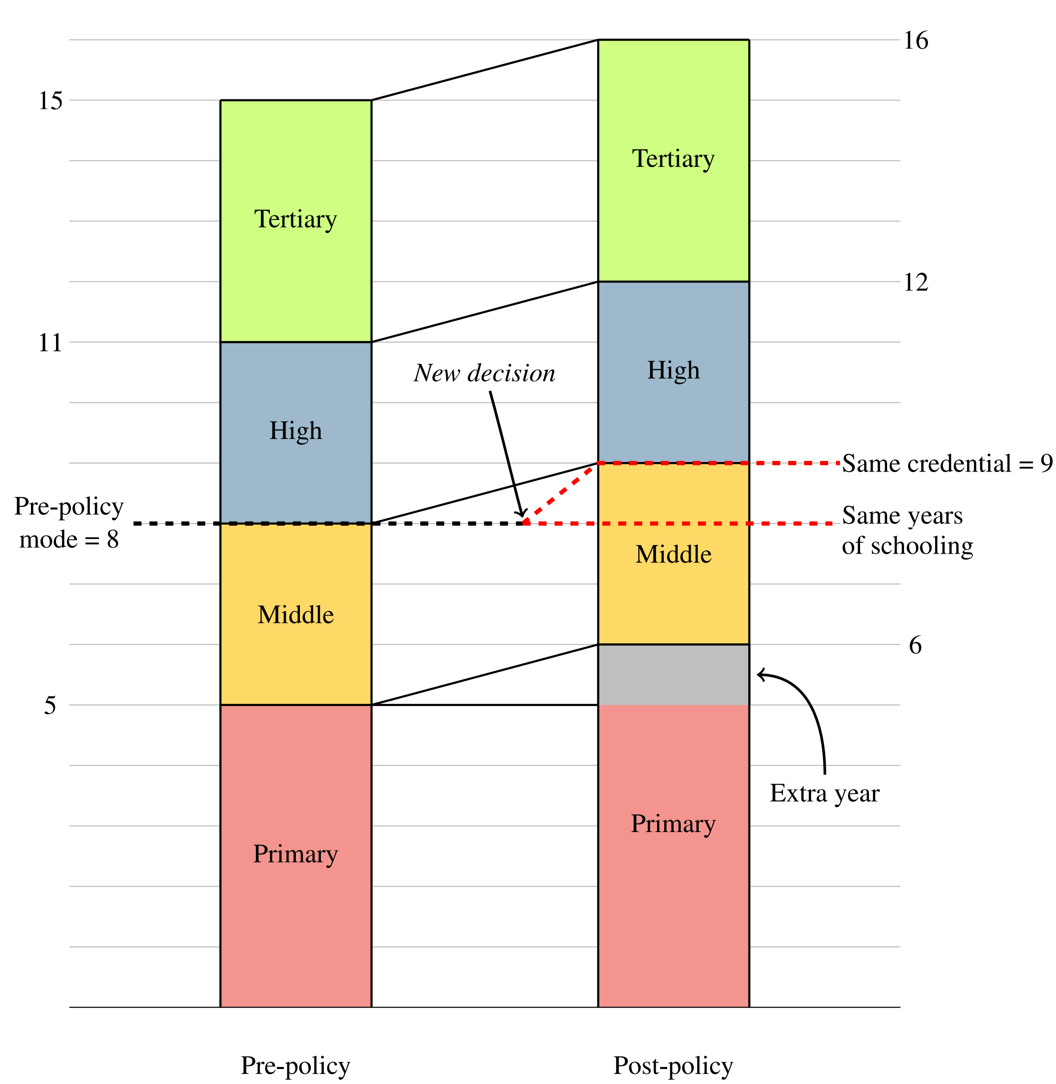

This figure depicts the length and structure of the Chinese education system before the policy is implemented in the left column, and after the policy is implemented in the right column. The y-axis represents the number of years needed to complete a credential under normal circumstances. Before the policy, it took five years to finish a primary school credential, as represented by the height of the box labeled "Primary" in the left column. Middle, High, and Tertiary refer to junior high school, senior high school, and university, respectively. To finish Middle School before the policy is implemented, an individual would need 8 years of school - five years of Primary and three of Middle, which can be read from the top line of the "Middle" box. The dashed bold line stemming from the y-axis at y = 8 indicates that prior to the policy, the median and mode of years of schooling is 8. The right column depicts the change introduced by the policy, adding an extra year to the time it takes to attain a Primary credential. The thick lines connecting the left and right columns depict the change in the total number of years it takes to earn each credential. The dashed red lines show the trade-off that the modal student faces after the policy.

We identify the causal effect of the policy on schooling and labor market outcomes using a regression discontinuity (RD) design with time as the running variable[3]. We compare outcomes of treated and untreated individuals within each locality, restricting our attention to those leaving primary school within a few years of when the policy took effect. This approach is similar to recent work studying the impact of a sudden change in compulsory education in the UK (Oreopoulos, 2006; Clark and Royer, 2013). We take advantage of the fact that the policy was rolled out gradually across China to flexibly capture differential regional trends, protecting against the risk of upward bias in our research design (Stephens and Yang, 2014). We determine if, when, and how the policy is implemented in each of China's prefectures by collecting and coding hundreds of official government documents, known as "educational gazetteers," which report implementation at the local level.

We first estimate the impact of this policy on years of schooling and the attainment of educational credentials. We show that total years of schooling increased by nearly one for affected individuals, and we find no evidence of individuals offsetting the extra year of primary school with a decrease in years of post-primary schooling or credential attainment. This pattern is driven by a high and inelastic demand for earning a credential. In data on completed years of schooling, we observe extensive bunching at credential attainment years prior to implementation of the policy. After the policy, we show that this bunching persists, and we find no evidence that the policy changed the characteristics of who earns which credential. This means that the vast majority of affected individuals chose to forgo a year of wages in the labor market to earn their final credential, and that the policy has reallocated nearly one trillion person-hours from the labor market to the pursuit of schooling to date[4].

We use these results to estimate a parameter of central interest in labor economics: the labor market returns to an extra year of schooling, holding highest educational credential constant. We exploit the large sample size of the 2005 Chinese mini-census to generate a precise but small estimate: the extra year increases monthly income by 2.03%. We find no evidence that the additional year affected other labor market indicators such as entrepreneurial activity, employment status, and type of employer, i.e., private vs. government, suggesting that the income gains we observe are likely to be flowing through the human capital accumulation channel and not through selection into different types of employment.

This result is quite robust. Its magnitude and precision are stable across specification and the choice of how many years of data around the first treatment year to include in the estimation sample. We conduct a permutation test, drawing 1,000 sets of placebo treatment years and using them to estimate placebo treatment effects. The distribution of these placebo estimates is centered around zero and the true estimate falls well beyond two standard deviations from the mean, strong

[3]This could also be called an event study, but we use the RD label, as we make heavy use of the machinery for establishing causal inference in that literature (Imbens and Lemieux, 2008; Lee and Card, 2008; Lee and Lemieux, 2010).

[4]China's educational yearbooks estimate that 332,321,868 children graduated with six years of primary school between 1984 and 2009. The number of students leaving primary school between 2010-2015 under the six-year regime, assuming negligible drop-out from primary school, is 90,045,459. The number graduating under this regime between 1981 and 1984 is not listed in the yearbooks. Using the proportions given in Figure A.1, we estimate that it is likely to be no more than a few million students. We assume individuals spend 2,000 hours working in a year. We multiply the number of affected individuals to date by the year of lost labor hours each forgoes to generate an estimate of approximately 0.85 trillion person-hours of labor reallocated to schooling between 1981 and 2015.

evidence that our estimate is not merely a mechanical result of our choice of research design.

In contrast to our estimate of a two percent return to the year of schooling, other credibly identified estimates from China and the developed world range from eight to 20 percent (Angrist and Krueger, 1991; Ashenfelter, Harmon, and Oosterbeek, 1999; Oreopoulos, 2006; Liu and Zhang, 2013). The naive correlation between years of schooling and earnings in our data is eight to 14 percent per year of schooling. We claim that much of the difference between our estimate and these much higher estimates is driven by our ability to shut down the signaling contribution to the returns to schooling by holding highest credential constant. We show three pieces of evidence in support of this claim.

The first is that the extra year appears to generate learning. We observe higher cognitive skills test scores among the treated group[5] as well as higher income gains among subgroups for whom we expect the extra year of learning to be particularly beneficial. Furthermore, the extra primary year is similar to other years of schooling in its curricular content, dispelling the concern that this is a "wasted" year. The second is a set of results showing that our estimate is unlikely to be biased downward by experience differentials[6], attenuation, general equilibrium effects, or a possible gap between our reduced-form treatment effect estimate and an IV estimate of the parameter. The third is that our estimate is similar in magnitude to that of Li, Liu, and Zhang (2012), the only other well-identified work we are aware of which attempts to disentangle the human capital and signaling contributions to the labor market returns to schooling in China. We use parameter estimates from our analysis and Li, Liu, and Zhang's work to make a back-of-the-envelope calculation of the relative importance of human capital accumulation and signaling in the naive labor market returns to schooling we observe in the cross-section. Our calculation suggests that between 57 and 60 percent of these returns are flowing through the signaling channel, with the remainder coming from the returns to human capital accumulation in school.

A primary impact of this policy is the reallocation of 850 billion person-hours from work to school. We perform a cost-benefit analysis to quantify the public finance implications of this reallocation. We borrow the framework used in Duflo (2001), estimating the lifetime value of the increase in monthly income conferred by the extra year of schooling and comparing it our estimate of the cost of the lost year of productive activity in the labor market for affected individuals over the lifetime of the policy. We generate four estimates, and in all but the most favorable scenario the costs of the policy exceed its benefits by at least tens of billions of US dollars.

Finally, we ask how generalizable are our findings that the demand for educational credentials is both high and inelastic with respect to the length of time it takes to earn these credentials? To do so, we examine Demographic and Health Surveys (DHS) data collected in 74 countries on completed years of schooling. We find patterns of bunching at credential attainment years similar to those we observe in China, and consistent with a high value of credential attainment, in 48 of these countries.

[5]This adds to recent evidence that there can be measurable adult returns to childhood interventions decades after the intervention runs its course (Heckman, 2006; Chetty et al., 2011).

[6]Our research design involves comparing treated individuals who have one more year of schooling and one less year of work experience to untreated individuals with one less year of schooling and one more year of work experience. We calculate Manski bounds to show that the wage premium to the extra year of experience that the untreated benefit from is low, on the order of 0-1.7 percentage points.

Our paper contributes new evidence to three lines of empirical research on the economics of education. The first line of research tries to disentangle the relative importance of signaling and human capital in driving schooling decisions and the returns to schooling (Lang and Kropp, 1986; Tyler, Murnane, and Willett, 2000; Bedard, 2001; Grenet, 2013; Clark and Martorell, 2014). The second line uses large changes in schooling policy from the developing world to assess the merits of different policy options and their distributional effects (Duflo, 2001; Banerjee et al., 2007; Lucas and Mbiti, 2012). The third shows the long-term impacts of early life schooling interventions on cognitive ability and earnings in adulthood (Chetty et al., 2011; Fredriksson, Öckert, and Oosterbeek, 2013).

The rest of the paper proceeds as follows. In Section II, we discuss the history of education in modern China and describe the policy we study. In Section III, we describe the data we use and our identification strategy. Section IV contains empirical results related to educational attainment and Section V provides our empirical results relating to the labor market. In Section VI we perform the cost-benefit analysis and draw policy implications. Section VII concludes.

II. A Brief History of Primary and Secondary Education in Modern China

China's education system has grown substantially in size and scope since the end of the Chinese Civil War in 1949. Education levels at that time were quite low: only 20% of the population was literate, and fewer than 40% of school-aged children were enrolled in school (Hannum, 1999). The new government increased spending on primary education after independence and vastly expanded the number of schools across the nation at all levels (Liu, 1993). During the period 1950-1980, a series of policy experiments and natural disasters racked China, leaving in its wake an educational system which varied greatly in scope and structure across provinces. Even so, literacy jumped to more than 50% by 1976 and average educational attainment rose to over 7 years (UNICEF, 1978). After the Cultural Revolution ended in 1976, China's education system moved towards standardization (Hannum et al., 2008). In January 1978, the *Full-Time Ten-Year Primary and Middle Education Teaching Plan (Draft)* mandated national harmonization of the length and structure of primary, middle, and high school[7] in all provinces. This set the length of primary school to be five years in schools across the country.

At the end of 1980, the Central Committee of the Communist Party of China and State Council issued the *Decision on Several Problems Relating to Universal Primary Education*, the policy whose changes we use for our analysis. This policy mandated that the total years of primary and secondary education be extended to twelve years, including a shift from five to six years of primary school[8]. It allowed gradual adoption of the primary school length change across localities, putting more pressure on urban schools (Liu, 1993). Figure A.1 plots national data on the proportion of students in six year (or equivalent) primary school systems, showing gradual adoption of the six year primary system between 1980 and 2005. About 60 percent of localities switched to a six year system between 1981 and 1993, relatively few made the change in the

[7] In China, middle school and high school are referred to as junior middle school and senior middle school. We refer to them here as middle school and high school to facilitate a layperson's understanding.

[8] In Shanghai and a few other localities, this policy was implemented instead by requiring that middle school last four years instead of three.

mid-1990's, and the rest shifted in the late 1990's and 2000's, reaching near-universal adoption in 2005. As explained in Footnote 4, we estimate that the policy has induced approximately 425 million students to spend a sixth year in primary school so far.

This policy was announced early on in Deng's time as China's de facto leader and at the beginning of the country's transition from a planned to a more market-oriented economy. One of Deng's early directives was that education should "face the demands of the new era and meet head-on the challenges of the technological revolution." The policy we study was implemented as part of his larger move in the late 1970's and early 1980's to prepare China's labor force to adapt to this new economic arrangement (Vogel, 2011).

This policy did not change the age at which children entered school, nor did it change the primary school, middle school, or high school curricula. Rather, the intent of the change was that primary school students be taught the same material over a longer period of time to ensure mastery of the curriculum.

We conducted a series of structured qualitative interviews with affected students, teachers, and parents to collect a richer account of what happened during this extra year and how it was perceived. Students and teachers reported that in the first year or two after the reform, the content of the extra year consisted of a review of what was covered in the fifth year of primary school and the addition of elective courses, such as physical education and music. After the adjustment process was completed, students and teachers reported that the primary school curriculum previously covered in five years was extended more smoothly over six years. In practice, this meant more time allowed for review and ensuring the foundational concepts of the primary curriculum were mastered by all students. Generally, respondents felt the extra year was most likely to have helped those of lower ability. More than half the respondents mentioned the loss of a year of productive work as the main downside of the policy.

The extra year of schooling posed logistical and personnel challenges. The policy required primary schools to hold and teach an extra cohort of children, but these schools were given no additional resources to do so. Gazetteer records and our interviews indicated that the additional burden of housing and schooling this extra cohort in a given primary school involved assigning more work to existing teachers and dividing up existing facilities, as opposed to building new structures and hiring new staff. When asked about the effect of this extra burden on the quality of education, respondents generally thought it unlikely to have a substantial impact. This claim is consistent with the rote nature of Chinese primary education during this time, which we argue is likely to dampen a possible negative relationship between class size and learning. As we document later in the paper, this additional burden was gradually offset by secularly declining cohort sizes over time.

The gazetteers document that the transition from five to six years of primary school was carried out in a number of ways. Table A.2 gives six examples of how the policy was enacted, taken from gazetteers in different implementing cities and counties across the country. In some cases, the transition was accomplished by enforcing the policy immediately, forcing all students who had not graduated from primary school, including those in fifth grade at the time, to remain in primary school an extra year. In other cases, it was accomplished by selecting a cohort of students (e.g. third graders) after which all students must complete six years of primary schooling. In other instances, a portion of the exiting cohort of students was sent on to middle school after their fifth year of primary school while the rest remained to finish a sixth year. This practice was explained in the gazetteers as a method to smooth the flow of students during the first year or two of transition, after which all subsequent cohorts would then take six years.

The decision of when to implement the policy was made at the local level. Though upper-level pressure certainly played a factor, as we discuss in Section III, most counties had the ultimate say on the year in which the switch was made[9]. Our data bears this out, and in section III.C we address the issues surrounding discretion in timing of implementation and the attendant concerns of omitted variable bias.

Our gazetteer data show no evidence of any other policy change which was regularly coincident with the change we study. A separate policy issued in April 1981 by the Ministry of Education mandated that the length of high school to be extended from two years to three by the end of 1985. This implementation occurred over a much shorter time frame than the extension of primary school from five to six years: by 1984, 90% of students in high school were in three year programs; in contrast, it was not until 2003 that more than 90% of primary school students were in six year programs (National Institute, 1984). In 1986, the Chinese government made middle school compulsory, but we show in Section VI that this law appears to have little impact on the middle school attainment of observations in our estimation sample. We argue this is due to two factors. One, there is widely documented porous enforcement of the law in rural areas (Fang et al., 2012). Two, in urban areas education levels are already high at the time of the policy announcement and so the law is binding for relatively few urban residents.

III. Data and Identification Strategy

This section describes the data sources and empirical methods of the paper. We show evidence that the main identifying assumptions for the research design are satisfied, and address a set of issues which could confound causal interpretation of our results.

III.A. Data Sources

Our sources of data are listed in Table A.1. There are two main sources of observational data: the 2005 Chinese mini-census and the 2010 wave of the China Family Panel Studies (CFPS). The 2005 Chinese mini-census collects basic data on family structure, highest educational credential attained, health, and income, and contains 2.6 million observations[10]. The CFPS is a new, nationally representative panel data set containing information from over 30,000 individuals in rural and urban China across 25 provinces, representative of 94.5% of China's population[11]. Summary

[9]Local educational gazetteers document that, in most cases, counties within a prefecture implemented the policy in the same year or within a few years of each other.

[10]Though the full sample is approximately 13 million observations, researchers are granted access to 20% sub-samples of the parent dataset.

[11]The data include all provinces but Tibet, Xinjiang, Inner Mongolia, Hainan, and Ningxia. The CFPS is conceived of as a panel, with six waves planned, taking place in 2010, 2012, 2014, 2016, 2018, and 2020. For this analysis, we use only the 2010 wave. The project is organized by a team of economists and sociologists at Peking University and collects a rich set of data on family structure, income, expectations, and several other social and economic indicators. Detailed

statistics on demographic, education, and employment characteristics of our sample population are given in Table A.3 for each data set, separately for rural and urban residents.

We also collect our own data from two sets of national archives to determine which observations in the census data were affected by the policy we study. The shift to six-year primary school was implemented at different times in different places both across and within China's provinces, as shown in Figure A.2. We hired a team of research assistants to read through county educational gazetteers stored in the Chinese National and Peking University Archives, to determine if, when, and how the policy was implemented in each locality[12]. Figure A.3 shows a page from one of these gazetteers.

We determine the year the policy was implemented, separately for rural and urban residents[13], in 280 of the 345 prefectures[14] in the census data. Of those 65 prefectures in the census we are unable to code, 45 either implemented the policy gradually across counties within a prefecture, so that we could not identify a unique prefecture-level treatment year, or instead changed to a system of five years of primary school with four years of middle school instead of the six primary plus three middle format we study. The remaining 20 had no record of implementing the policy in the currently available educational gazetteers.

Determining Treatment Status in the CFPS Data

We use the gazetteer data to identify which individuals in the census are treated by the policy and which are not. In the CFPS dataset, the location of observations is anonymized to the provincial level, which prevents us from using the gazetteers to determine treatment status. Nonetheless, the CFPS has a few traits which make it particularly desirable. It collects detailed data on how many years individuals spend in each level of schooling and identifies which individuals reside in a given county. The gazetteers document that the policy is sometimes implemented at different times between counties within prefectures. As a result, analysis at the county level is important to understand precisely how the policy was rolled out in each location.

We apply a mean-shift algorithm (Fukunaga and Hostetler, 1975) to the CFPS data to determine treatment status for observations in the CFPS. Our algorithm generates, for each county, the most likely cohort in which the number of years spent in primary school jumps from five to six[15]. Its implementation in our context is straightforward - for observations in a given county, we regress individual-level years of primary school on a constant and an indicator function for having graduated in or after a given year:

$$s_i = \gamma_0 + \gamma_1 \cdot \mathbf{1}\{t_i \geq t^*\} + \varepsilon_i \tag{1}$$

We estimate 27 regressions for each county, corresponding to every possible treatment year in our data, $t^* \in [1981, 2007]$. In this equation, s_i is the number of years individual i spent in primary

information about the sampling structure and overall plan for CFPS is available in Lv and Xie (2012).

[12]Recent work by Almond, Li, and Zhang (2013) uses Chinese gazetteers to identify when land-reform policy was implemented in different counties across the country.

[13]In many prefectures, urban implementation occurred before rural implementation.

[14]Counties, prefectures, and provinces are the Chinese geographic divisions of interest to this study. There are several counties in a prefecture and several prefectures in a province.

[15]This mean-shift approach is similar to that used Munshi and Rosenzweig (2013).

school, t_i is the year in which she graduated from primary school, and ε_i is an i.i.d. error term. The year (t^*) with smallest sum of squared residuals (ssr) is the predicted treatment year for that county[16]. This exercise generates a treatment year for each county in our estimation sample[17]. In Appendix 3, we use national statistics and the application of both archival and mean shift methods to a third observational data set, the China Labor-Force Dynamics Survey[18], to corroborate the reliability of the mean shift method's identified treatment years.

III.B. Empirical Strategy

Our identification strategy is a simple regression discontinuity design with a discrete running variable: distance in years between an individual's birth cohort and the first cohort affected by the policy change we study (Lee and Card, 2008). We compare outcomes of individuals finishing primary school just before the policy is implemented in a given locality (county or prefecture) to those in the same locality finishing primary school just after implementation. The gradual rollout of the policy across time and space allows us to make this comparison while controlling flexibly for cohort, place, and cohort-by-region fixed effects.

For causal interpretation of our results, we require that within our geographical unit of interest, there is continuity in the conditional expectation of the outcome variable across the assignment threshold (Lee and Lemieux, 2010). We test this assumption on three fundamental predetermined characteristics which could affect our dependent variables, either through sorting or another selection mechanism: relative cohort size, gender composition of cohort, and proportion of individuals with a household registration certificate (hukou) from an urban area. Figure II plots these data, condensed to distance-to-treatment-year (the number of years between an observation's cohort and the first affected cohort) bin means. This figure shows no visible discontinuity at the treatment threshold. Due to the discrete nature of the running variable, we cannot run a McCrary (2008) test for bunching. Instead, as recommended in Lee and Card (2008), we use our main regression equation to estimate the "effect" of the treatment on the three predetermined variables for each dataset. In all cases we fail to reject a zero effect.

Following Imbens and Lemieux (2008) and Lee and Lemieux (2010), our main estimating equation is an ordinary least squares regression of y_{lci}, the outcome of interest for individual i in birth cohort c and locality[19] l, on a short set of key regressors:

$$y_{lci} = \beta_0 + \beta_1 * Treated_{lc} + \beta_2(t_{lc}|t_{lc} \geq 0) + \beta_3(t_{lc}|t_{lc} < 0) + \beta_4 V_i + \lambda_c + \mu_l + \eta_{cr} + \varepsilon_{lci} \qquad (2)$$

[16]An example of this process is shown in Figure A.4, which shows the histogram of cohort-mean years of primary school in a county and plots the ssr estimates generated by Equation 1 for each treatment year. The ssr sequence reaches its minimum at 1997, where we also observe a clear shift upwards in mean years of primary school from approximately five to six.

[17]Beginning with 162 counties in the CFPS, we exclude the 18 counties from Shanghai, as they implemented the policy by extending middle school instead of primary school. Of the remaining 144, we only those 112 counties in which we can detect a clear policy change. Appendix 3 lists the inclusion criteria used to determine this sample. Our empirical results are robust to using data from all 144 non-Shanghai counties.

[18]The China Labor-Force Dynamics Survey (CLDS) is a panel survey similar to CFPS. We use it to corroborate the reliability of the Mean Shift method we use on CFPS. We do not use the CLDS for our main regressions as it has neither the large sample size of the census nor the fine-grain locality information of the CFPS (CLDS indicates only which prefecture individuals are in, not which county).

[19]As mentioned earlier, this is at the county-level in the CFPS data and prefecture in the census.

Figure II: Predetermined characteristics across the treatment threshold

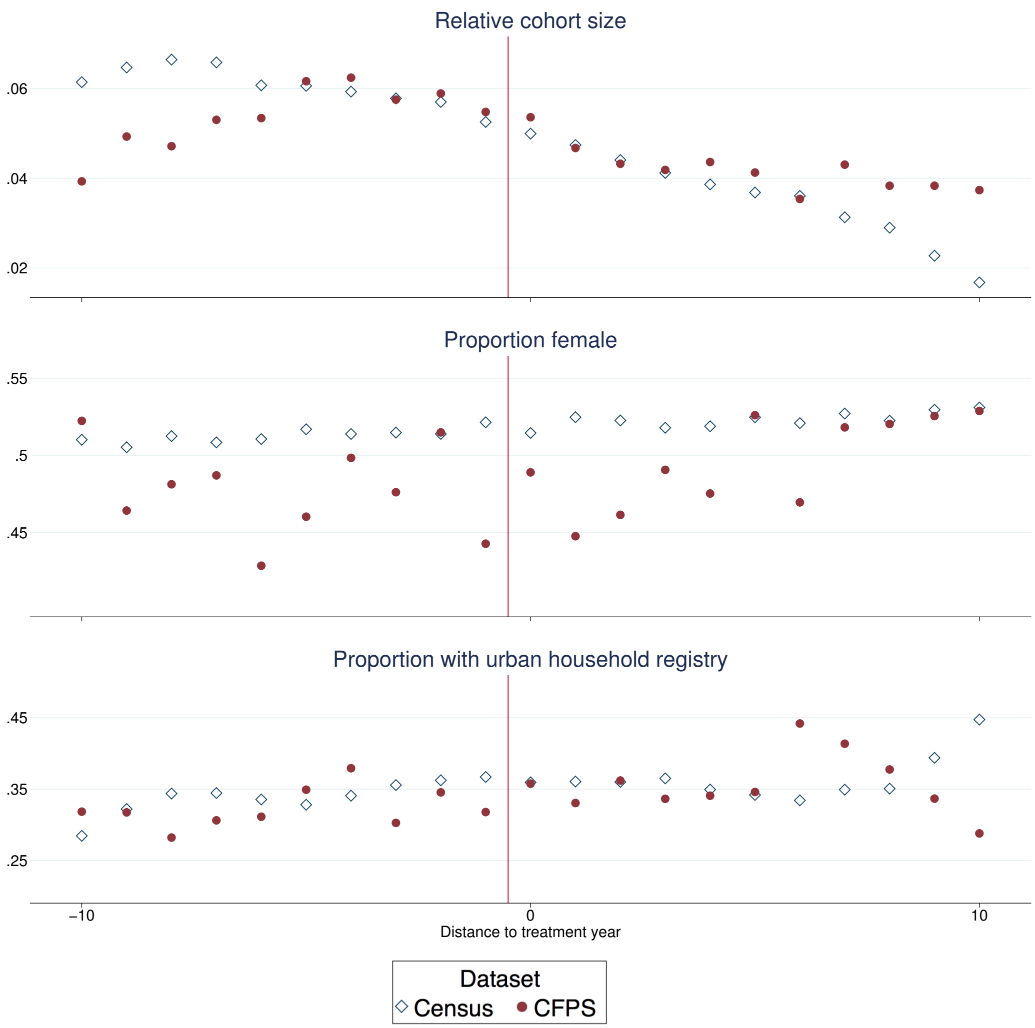

This figure shows suggestive data testing our main exclusion restriction, i.e., that there is continuity in the conditional expectation of the outcome variable across the treatment threshold. The red vertical line demarcates the untreated group (to its left) from the treated (to the right). The units on the y axis are indicated in the title above the sub-figure.

Here $Treated_{lc}$ is an indicator variable equal to 1 if the individual belongs to a cohort finishing primary school in or after the first affected cohort in her locality. t_{lc} is the locality-specific distance-to-treatment-year. We estimate the coefficient on distance-to-treatment-year separately for treated and untreated groups to account for pre- or post-policy trends (e.g., the possibility that the effect may differ over time elapsed since the treatment year, as counties get better at implementing the policy) in order to ensure that β_1 captures only the difference between pre- and post-policy means[20] (Gelman and Imbens, 2014). V_i is a vector of predetermined characteristics which includes, at the individual level, gender, ethnicity, residence permit status, and urban/rural residence, which can vary within a county or prefecture. Locality (μ_l) and either cohort (λ_c) or cohort-by-region[21] (η_{cr}) fixed effects are also included in all specifications unless otherwise stated.

Following Lee and Card (2008), we test that our estimated coefficients are stable across the choice of how many years around the treatment threshold we include in the estimation sample. We show this stability for two main empirical results. For the rest, our estimates use the sample limited to five years before or after the first treated cohort in each place. All regression results we present use robust standard errors clustered at the county or prefecture level (Bertrand, Duflo, and Mullainathan, 2004). We restrict our time frame to cohorts leaving primary school between 1976[22] and several years before the sample is drawn (1995 in the census data and 2003 in the CFPS data) to give most individuals enough time to finish their schooling career before being observed.

III.C. Potential confounders

The implementation of this policy across space and time was decided upon by local (province, prefecture, and county-level) bureaucrats. We have shown evidence that our main identification assumption is upheld and, as we are comparing within prefectures and counties, we do not need the timing of the policy to be randomly assigned across localities (Black, Devereux, and Salvanes, 2005; Meghir and Palme, 2005). If, however, there were another policy or external phenomenon correlated with both when the policy was implemented in a given locality and the later schooling decisions or labor market outcomes of affected individuals, our results would suffer from omitted variable bias.

Though it is impossible to rule out the existence of any external driver of both the timing of the policy and our outcomes of interest, we fail to find evidence of such a factor among a set of likely candidates. First, as recommended by Stephens and Yang (2014), we control for cohort-by-region fixed effects. If such an external phenomenon were geographically auto-correlated, for example due to intra-provincial policy coordination, these fixed effects would dampen its impact on our estimates. Second, we show the geography of the timing of implementation in each of China's prefectures according to archival records. Figure A.2 provides a heat map of prefecture

[20]In Table A.2 we give six examples from gazetteers of how the policy is implemented which speak to the need to control for the possibility of implementation varying over time.

[21]China can be separated into four well-recognized regions: East, Northeast, Central, and West.

[22]This coincides with the end of the Cultural Revolution and the end of the chaos it brought to the educational system of China.

implementation years, with lighter shades indicating earlier implementation. This map shows a wide distribution of timing of implementation with no obvious geographical pattern aside from later implementation in some prefectures in the central region. Third, we searched the gazetteers for mention of a policy or external influence that was regularly coincident with the implementation of the policy we study and found no such pattern. As these documents serve as official records of policy implementation, this strongly suggests absence of a consistent confounder.

Finally, as in Black, Devereux, and Salvanes (2005), we use OLS to estimate the relationship between timing of policy implementation and prefecture-level characteristics (e.g. gender composition, mean income, and proportion of respondents working for the government) for those prefectures in our estimation sample. While correlations here do not necessarily threaten causal interpretation of our results, they give a descriptive account of what may have driven implementation choice. The results of this regression are given in Table A.4, and show that only the central and northeastern regional dummies are statistically significant correlates of timing of implementation, consistent with the map in Figure A.2.

A final concern is the potential for migration to bias our results. We cannot reliably estimate treatment effects for migrants for two reasons: there are far fewer of them in our data and their treatment status is more difficult to pin down because we do not have data on precisely when they moved. In light of these limitations, we exclude migrants from our analyses. If the treatment effect is different for migrants and non-migrants or the policy affects who is likely to migrate, our estimates will differ from the true population average treatment effect. As there are far more rural-urban than urban-urban migrants, this is more likely to be a concern in rural areas. We test for an effect of the treatment on cohort size and characteristics of individuals in our rural sample, and find no significant relationship between treatment status and cohort size or gender composition. This suggests the extra year of primary school did not affect the propensity to migrate for the population as a whole, or differentially for men and women. Beyond these tests, there is little we can do about this concern, but the sample size of the census gives us the statistical power to conduct these tests with precision and the relatively high response rate of the CFPS (97% for households, 72% for identified adults within households) suggests that, at the very least, any migration-induced selection bias will be minimal (Lv and Xie, 2012).

IV. Empirical Results - Schooling

In this section, we estimate the impact of the policy on individuals' schooling outcomes. First, we show that the policy was indeed effective at extending the number of years individuals spent in primary school by one year. We then estimate the impact of this change on individuals' later schooling outcomes. These outcomes include years spent in post-primary schooling, whether or not an individual attains one of two post-primary credentials (a middle school or high school diploma), and drop-out. We finish this section looking at the effect of the policy on vulnerable subgroups and on the characteristics of individuals with each credential.

We first examine whether the policy had its desired effect of increasing primary school for affected individuals. Figure III plots distance-to-treatment-year bin means of the proportion of individuals spending six or more years in primary school in our CFPS sample, overlaid on es-

Table I: Effects of the policy on schooling outcomes - average treatment effects

Outcome	CFPS	2005 Census
Probability: at least six years of primary school	.547*** (.0294)	
Years spent in all levels of school	.66*** (.209)	
Years of post-primary schooling	.0928 (.204)	
Highest credential: at least middle school	.00368 (.0319)	-.00495 (.00303)
Highest credential: at least high school	.000299 (.0263)	.00628* (.00327)
Dropped out of school, any level		-.00017 (.00137)
Observations in sample	2,680	243,548

* p<0.10, ** p<0.05, *** p<0.01. Each cell presents a treatment effect estimate from a separate regression with the relevant robust standard error below, in parentheses. Standard errors are clustered at the county (CFPS) or prefecture (census) level.

timates of their confidence intervals. Prior to implementation of the policy, the proportion of students spending at least six years in primary school is between 20 and 30% of the population. This 20-30% comprises mainly individuals performing poorly in school who are made to repeat a grade[23]. At the policy implementation year it jumps to over 80%, increasing to nearly 100% over the next ten years. Results from the regression analog to this exercise are presented in the first row of Table I. We estimate that the treatment causes a 0.547 increase (SE 0.029) in the probability of taking a sixth year of primary school for those who graduate from primary school within five years after the policy is implemented.

This is an underestimate of the "first stage" of the policy: a sixth year of primary schooling post-policy is a deliberate expansion of the primary curriculum, as opposed to a forced repetition of the fifth year of primary school. Furthermore, as shown in Figure III, this probability estimate increases as we increase the bandwidth around the treatment year. We argue these circumstances give us a sharp rather than a fuzzy discontinuity, allowing us to consider none of the pre-policy group and nearly all of the post-policy group to be "treated."

Recall that this policy was implemented in each locality at a time when over 75% of students went on to get at least some post-primary schooling. Affected individuals could potentially hold

[23]Data from the baseline wave of the China Education Panel Survey (CEPS), a new, nationally representative dataset collected by scholars at Renmin University of China, corroborate this claim. In the CEPS dataset, approximately 16% of surveyed individuals repeated at least one year in school.

Figure III: Proportion of students spending at least six years in primary school before and after policy change

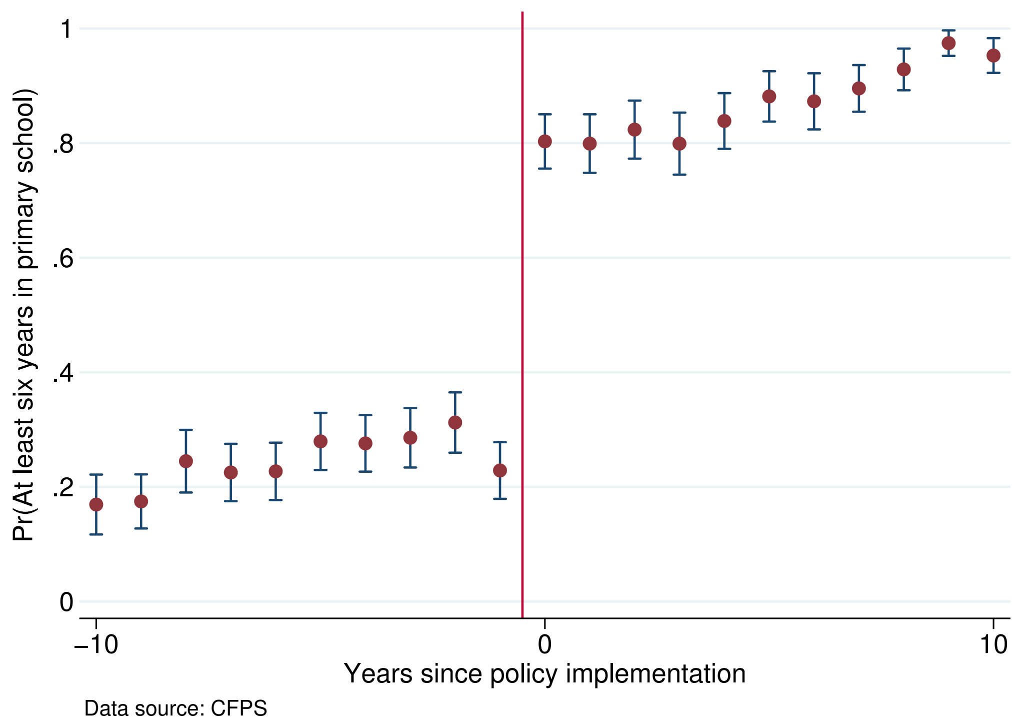

This plot shows distance-to-treatment-year bin means of the proportion of individuals spending at least six years in primary schooling before and after the policy. The vertical line separates the affected (to the right of the line) and unaffected (to the left) cohorts.

Figure IV: Distribution of post-primary schooling before and after policy change

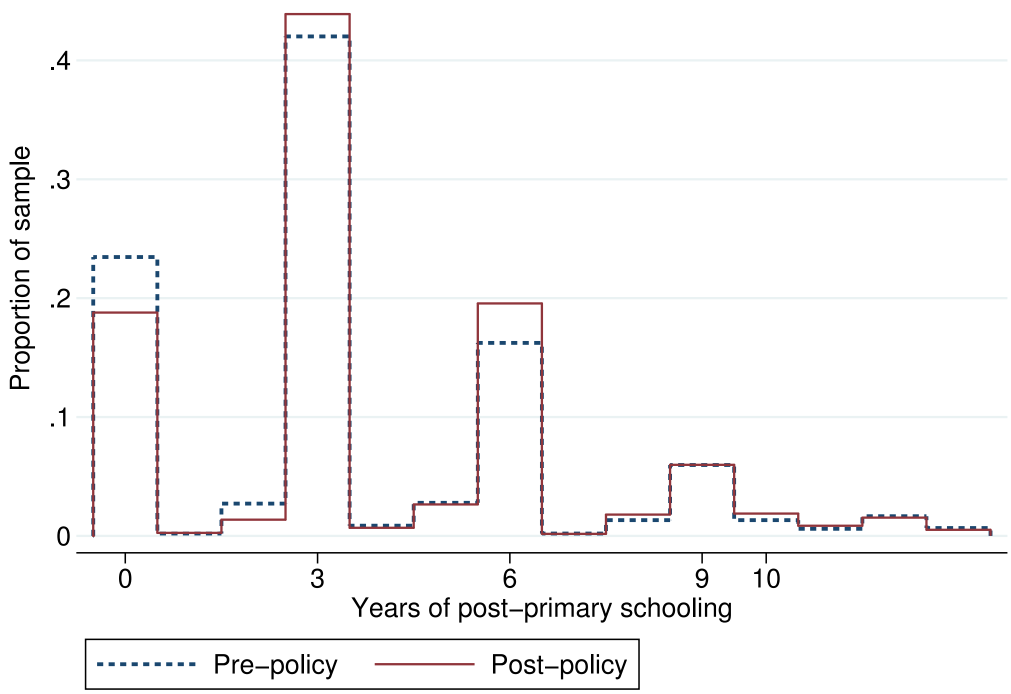

This plot shows the probability mass functions of post primary schooling for observations in the CFPS data graduating from primary school within five years before (pre-) and after (post-) the policy implementation year in each county. Note that the bunches at 0, 3, 6, 9, and 10 correspond to primary, middle school, high school, technical college, and university credential attainment years, respectively.

total years of schooling constant by offsetting the additional year of primary school with one less year of post-primary school[24]. Figure IV shows, separately for untreated and treated observations finishing primary school within five years of the treatment year, the distribution of post-primary schooling. This figure summarizes our main empirical results related to the effect of the policy on post-primary schooling. We see extensive bunching at credential attainment years and little visible difference between the treated and untreated groups in either the location or the magnitude of this bunching[25].

The regression results for our schooling outcomes are given in the rest of Table I. The second row shows our estimate of the effect of the policy on total years of schooling to be 0.660, significant at the 1% level. This implies that the vast majority of Chinese citizens induced by the policy to attend an extra year of primary school chose not to offset this with less post-primary schooling.

[24]This assumes the extra year of primary school does not confer a major gain in skills sufficient to induce individuals to proceed further in schooling. We provide evidence for this assumption in the next section.

[25]Zero years of schooling is the end of primary school, three is the end of middle school, six the end of high school, nine the end of technical college, and ten the end of university.

Our estimate of the effect of the policy on years of post-primary schooling is positive (0.093) but statistically indistinguishable from zero. The standard errors we generate can exclude anything larger than a 0.32 year decrease in post-primary schooling in response to the extension of primary school and also admit positive estimates of up to a 0.5 year increase. In Panel A of Figure V, we show that our point estimate of the impact of the policy on post-primary schooling is stable over nine different bandwidth choices, as recommended in Lee and Card (2008). In no case are we able to reject a zero effect[26].

We next use the census data to examine the effect of the policy on credential attainment. The census has coarser data on educational achievement (only highest credential attained, not years spent in each level of schooling) but is two orders of magnitude larger than the CFPS data. In the fourth and fifth rows of Table I, we estimate the effect of the policy on whether or not an individual earns at least a middle school credential and whether or not she earns at least a high school credential, using both census and CFPS data. The effect of the policy on middle school completion, estimated using the census data, is negative but small and insignificant (0.0049, SE 0.0030, from an untreated group average of 0.725). We can reject anything larger in magnitude than a one percentage point decrease on the probability of completing middle school. The estimated effect on finishing high school is small, positive (0.0063), and significant at the 10% level. We find no effect on the probability of dropping out. These small standard errors speak to the statistical power the census affords us relative to the CFPS in measuring even small effects. In short, we find no evidence of a decrease in post-primary schooling to offset the lengthening of primary school.

One possible explanation for this overall pattern of no net change in post-primary schooling is a change in composition of who earns which credential. The zero effect could mask two countervailing phenomena: first, some individuals advancing further than they would by virtue of the skills gained in the extra year, and second, others reducing post-primary schooling by an entire credential. To test for this possibility, we perform two exercises. First, we look for changes among those subgroups we would expect to be most likely to be induced by the policy to offset the extra primary year with fewer post-primary years; second, we explicitly test for changes in composition of background characteristics at each level of schooling.

Previous work has shown that Chinese households in the 1980's and 1990's often chose to allocate fewer resources to women (Li, 2003). Income is also much lower in rural areas of China than in urban areas, and women from rural areas are thus doubly disadvantaged. We anticipate the extra year of forgone wages needed to earn a credential poses the greatest burden for these three groups. If this is true, we are most likely to find a downward shift in their post-primary schooling to offset the extra year of primary school.

We test these predictions in Table II, which shows the subgroup-specific treatment effect estimates for the same outcomes examined in Table I. The estimated treatment effects are largely negative, as predicted, and consistently so for rural women, the most disadvantaged group (Con-

[26]In Figure A.5, we show the event study graphs for these effects for both the raw data (left column) and cohort-demeaned residuals (right column) for total years of schooling (top row) and years of post-primary schooling (bottom row). We see the same patterns as in the regression coefficients: an upward jump of about one year of total schooling at the treatment threshold, and no downward jump in post-primary schooling at the treatment threshold, both for the raw and demeaned data.

Figure V: Stability of main regression estimates across bandwidth choice

Panel A - Treatment effect estimates and confidence intervals for post-primary schooling

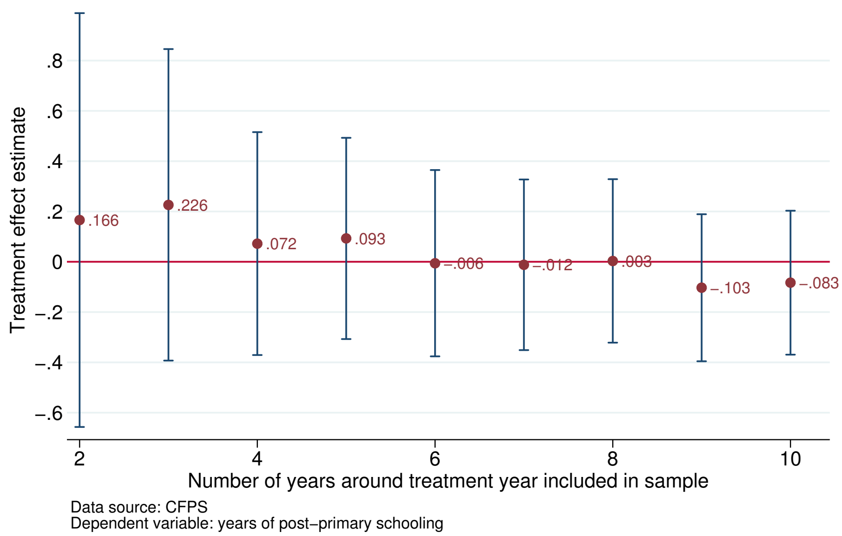

Panel B - Treatment effect estimates and confidence intervals for log monthly income

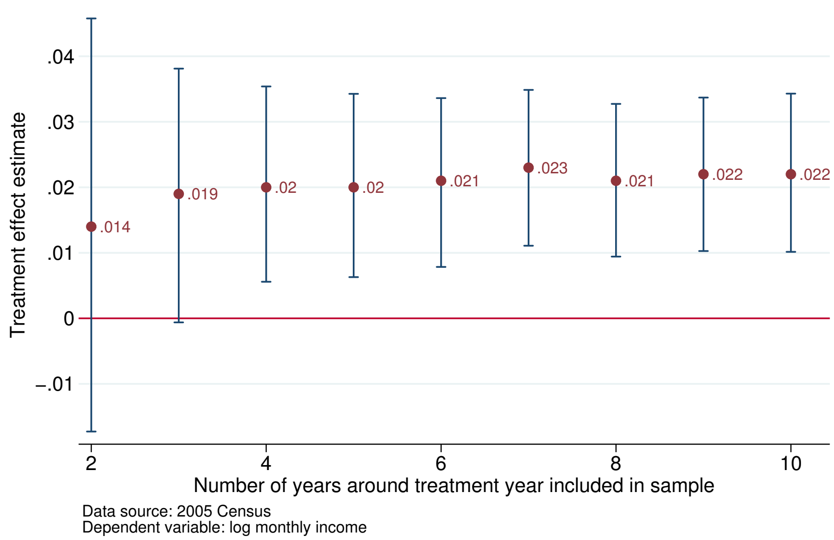

Table II: Effects of the policy on schooling outcomes for vulnerable subgroups

Outcome	Female	Rural	Rural female
Years of post-primary schooling			
CFPS	-0.2020	0.2357	-0.0472
	(0.2951)	(0.2673)	(0.4008)
Graduated from middle school			
CFPS	-0.0355	0.0550	0.0304
	(0.0498)	(0.0462)	(0.0696)
Census	-0.0057	-0.0049	-0.0044
	(0.0043)	(0.0046)	(0.0063)
Graduated from high school			
CFPS	-0.0221	0.0238	0.0133
	(0.0400)	(0.0195)	(0.0352)
Census	0.0043	-0.0016	-0.0030
	(0.0039)	(0.0025)	(0.0028)
Dropped out of school, any level			
Census	0.0037*	-0.0003	0.0047*
	(0.0019)	(0.0021)	(0.0028)
Observations in CFPS sample	1,277	1,293	592
Observations in Census sample	126,081	157,308	81,490

* p<0.10, ** p<0.05, *** p<0.01. Each cell presents a treatment effect estimate from a separate regression with the relevant robust standard error below, in parentheses. Standard errors are clustered at the county (CFPS) or prefecture (census) level.

nelly and Zheng, 2003). The magnitude of the estimates is uniformly small, however, and only for dropout rates do they reach statistical significance at the 10% level. We estimate but do not present effect estimates for other groups (such as men, those from western and non-western provinces, and urban areas), which we find to be more consistently positive but similarly small relative to their respective standard error estimates.

The second exercise estimates a version of our main empirical specification to test for compositional changes. We replace the single treatment variable with four dummy variables for the treatment interacted with an individual's highest educational credential (primary, middle, high, or tertiary). For outcome variables, we use a set of predetermined characteristics to proxy for household resources allocated to the child and scholastic ability, the most likely predictors of adjustment on the credential margin. We proxy for resources allocated to the child with number of siblings and gender. We use parents' highest credential (mother's and father's separately) to proxy for scholastic ability. Wald tests of the equality of the treatment-by-credential level coefficients tell us whether the proportion of individuals with the predetermined characteristic of interest holding a given credential changes, relative to that proportion among those holding other credentials, across the treatment threshold. We use the CFPS data for these tests, and fail to reject equality of the treatment-by-credential coefficients on number of siblings, mother's and father's highest educational credential, and gender (p-values 0.740, 0.665, 0.135, and 0.660, respectively). We conclude from these analyses that the characteristics of who earns which credential are unlikely to change substantially as a result of the policy.

V. Empirical Results - The Labor Market

In the previous section we studied the effect of a policy which adds an extra year to primary school on schooling outcomes, assuming individuals could adjust their level of post-primary schooling up or down to compensate for or reinforce the policy's effects. We found that the policy induced the vast majority of Chinese citizens to spend an extra year in primary school while leaving their highest educational credential unaffected. In this section, we will treat the variation from this policy as an experiment which induced all affected individuals to spend an extra year in school while holding their highest credential constant. In the first subsection, we estimate the effect of the policy on labor market outcomes and run robustness checks. In the second, we show several pieces of evidence which suggest that our treatment effect estimate is quite close to the parameter we are after, isolating the labor market returns to the human capital accumulated during a year of schooling from the signaling effect of receiving an additional credential that often confounds such work (Weiss, 1995; Card, 2001).

V.A. Estimation results and robustness

We use Equation 2 to estimate the effect of this additional year of schooling on various labor market outcomes, including employment status, type of employment for the employed (i.e., self-employment, government sector employment, and private sector employment), and monthly income. Though China was strictly a command economy as recently as 1978, reforms enacted in the 1980's and 1990's pushed the Chinese labor market to more closely resemble that of a market economy as early as the late 1990's (Cai, Park, and Zhao, 2008).

We use the 2005 mini-census data for all of the analyses in this section due to its large sample size[27]. We restrict our attention to urban residents, as in rural areas treatment effect estimates would be muddied by the impact of the treatment on the decision to work in agriculture or not and selective loss to migration is more of a concern. Our main dependent variable of interest is the natural logarithm of monthly income[28]. We also investigate the effects of the policy on whether the individual is employed and whether she is employed in a government job, in the private sector, or works as an entrepreneur. In the regression results presented in this section, we add highest credential fixed effects to the right hand side of Equation 2 and use the same sample restricted to five years on either side of the treatment year for estimation[29].

[27]By 2005, we expect most workers to earn wages that are at least strongly correlated with their relative productivity (Zhang et al., 2005).

[28]When estimating the effect of the policy on income, we drop those 266 observations from the five-year bandwidth estimation sample (163 in the treatment group, 103 in the control; out of 66,691 observations) who are working but report zero monthly income.

[29]We do not have labor histories for individuals, and so cannot apply the normal Mincerian equation using years worked (experience) as an independent variable. Instead, we assume individuals are gaining experience in each year they are not in school. Under this assumption, the birth year (cohort) and credential level fixed effects are a sufficient statistic for the experience of the individual.

Regression results are given in Table III. We find no evidence that the policy had any effect on whether or not an individual is working, with a treatment effect very close to zero (0.26 percent, from a treated group mean of 77.7 percent) and standard errors precise enough to reject anything larger than a 1.2 percentage point increase or a 0.8 percentage point decrease in this probability. Our estimate of the effect of the policy on whether the individual works for the government (as opposed to for the private sector or as an entrepreneur) is similarly precisely estimated and indistinguishable from zero. This result suggests that the policy is unlikely to have had a large effect on whether or not an individual is working and, if so, whom she is working for.

Next, we estimate the effect of the policy on the natural log of monthly income. We add employer-type fixed effects to this specification to more precisely estimate our parameter of interest, the income returns to the extra year of schooling[30]. Our first specification uses cohort and locality (prefecture) fixed effects. Here we find a gain of 1.94% in monthly income, statistically significant at a 99% confidence level. Recent work (Stephens and Yang, 2014) shows that previous efforts to estimate the returns to schooling using strategies similar to ours may have been biased upward and suggests including cohort-by-region fixed effects to mitigate this bias. Implementing this recommended approach, we next estimate Equation 2 with the addition of cohort fixed effects specific to each of China's regions (East, Northeast, Central, and West). Our treatment effect estimate increases slightly, to a 2.03% gain, and remains significant (the standard errors change by less than three hundredths of a percentage point). Panel B of Figure A.5 shows the corresponding event study plot.

We next explore heterogeneity in treatment by subgroups, shown in Table IV. The coefficients shown here are from an estimating equation similar to equation 2, only replacing the single treatment variable with interactions between the treatment variable and a dummy for membership in the mutually exclusive and exhaustive subgroups of interest (e.g. men and women, government and non-government workers) and excluding the un-interacted treatment variable from the equation for ease of interpreting each subgroup-by-treatment coefficient. Panel A shows that the estimated gain from the additional year is monotonically decreasing in highest educational credential, consistent with the goals of the policy and with the proportional contribution of an additional year of schooling diminishing as total years of schooling increases. As mentioned in the introduction, this result is also consistent with a study that identifies increased instructional time as a key mechanism contributing to the success of a set of New York City charter schools in raising achievement among underprivileged students (Dobbie and Fryer, 2013).

Panel B shows a higher return to the extra year for women than for men, consistent with other work on returns to schooling by gender in urban China (Hannum, Zhang, and Wang, 2013). Panel C shows that private sector workers appear to enjoy all of the wage premia from the extra year. This difference is unsurprising, as pay is almost certainly more closely linked to the relative productivity of labor in the private sector than in the government (Li et al., 2012). Independently run Wald Tests reject equality of the coefficients for both differences.

We attempted to investigate the possibility of further conditional treatment effect heterogeneity, e.g. gender-by-occupation or gender-by-education level returns, but our research design is

[30]Results omitting these fixed effects are similar in magnitude, varying by less than 0.5%.

Table III: Effects of the policy on labor market outcomes

Outcome	Estimates
Panel A: average treatment effects	
Currently employed	0.0026
	(0.0052)
Works for government	0.0035
	(0.0062)
Log of monthly income, using cohort and place fixed effects	0.0194***
	(0.0073)
Log of monthly income using cohort, place, and cohort-by-region fixed effects	0.0203***
	(0.0071)
Panel B: effect on log of monthly income, by highest credential†	
Primary school	0.0694***
	(0.0279)
Middle school	0.0503***
	(0.0095)
High school	0.0257***
	(0.0092)
Tertiary schooling	-0.0098
	(0.0093)
Panel C: effect on log of monthly income, by gender	
Men	0.0081
	(0.0085)
Women	0.0375***
	(0.0086)
Panel D: effect on log of monthly income, by employer	
Government	-0.0139
	(0.0086)
Enterprise	0.0327***
	(0.0076)

* p<0.01, ** p<0.05, *** p<0.01. Data: census. For first row, N=86,240. For all other regressions, N=66,425. Panel A presents results from four different regressions. In panels B, C, and D, we present coefficients from a single regression as specified in the panel title. Panels B, C, and D present coefficients of a dummy variable for membership in the group given in the left column (e.g. those whose highest credential is primary school) interacted with the treatment dummy. All samples include only urban residents and non-migrants. Robust standard errors are given below the coefficient estimate in parentheses and are clustered at prefecture level. †Panel B coefficients, when weighted by proportion of sample given in column 4 of Table V, sum to 0.0213.

Table IV: Effect estimates on log monthly income, by subgroup

	Dependent variable: log monthly income
Panel A: Effect estimate by highest credential[†]	
Primary school	0.0694***
	(0.0279)
Middle school	0.0503***
	(0.0095)
High school	0.0257***
	(0.0092)
Tertiary schooling	-0.0098
	(0.0093)
Panel B: Effect estimate, by gender	
Men	0.0081
	(0.0085)
Women	0.0375***
	(0.0086)
Panel C: Effect estimate, by employer	
Government	-0.0139
	(0.0086)
Enterprise	0.0327***
	(0.0076)

*p<0.10, ** p<0.05, *** p<0.01. Data: census. N=66,425. Each panel presents result from a separate regression. Results presented are estimates of the coefficient of a variable which is a dummy for membership in the group given in the left column (e.g. those whose highest credential is primary school) interacted with treatment status. All samples include only urban residents and non-migrants. Robust standard errors are given below the coefficient estimate in parentheses and are clustered at prefecture level. [†]Panel A coefficients, when weighted by proportion of sample given in column 4 of Table V, sum to 0.0213.

too data-intensive to precisely estimate these effects, even using the census data. Comparing the treated and untreated, within subgroups of subgroups in each locality, limited to a narrow bandwidth around the treatment year, leaves us with too few observations per locality to generate precise estimates using the RD design as specified.

Next we perform a few robustness checks. In Panel B of Figure V, we present our estimates of the effect of the policy on log monthly income and their confidence intervals for the same nine different choices of sample bandwidth as in Panel A of that figure. Both the magnitude of the coefficient and its ability to reject a zero effect are stable across a wide range of bandwidth choice, suggesting that our result is robust to large changes in the estimation sample. Table III shows that the coefficient estimate is also robust to the inclusion/removal of cohort-by-region fixed effects.

There is still concern that our research design, if specified incorrectly, could mechanically generate a difference between the treated and control groups unrelated to the effect of the policy. To test for this, we conduct a permutation test. This is a Monte Carlo-type analysis which estimates the distribution of treatment effects for years in which the policy does not occur. First, we draw 1,000 placebo treatment years for each prefecture (sampled from the full support of the estimation sample's potential years, 1981-1997). Then, using the treatment status assigned by these placebo years, we estimate the placebo treatment effect on wages for each draw. Figure VI gives the probability density function for these estimates. The placebo treatment effect estimates are normally distributed, with a mean of -0.0010 and a standard error of 0.0084, putting the true estimate of 0.0203 well beyond two standard deviations from the mean. We conclude that the sign and significance of our estimates are not merely a mechanical result of our research design.

Figure VI: Permutation test results - distribution of effect estimates of placebo treatment on log monthly income from 1,000 draws of placebo years

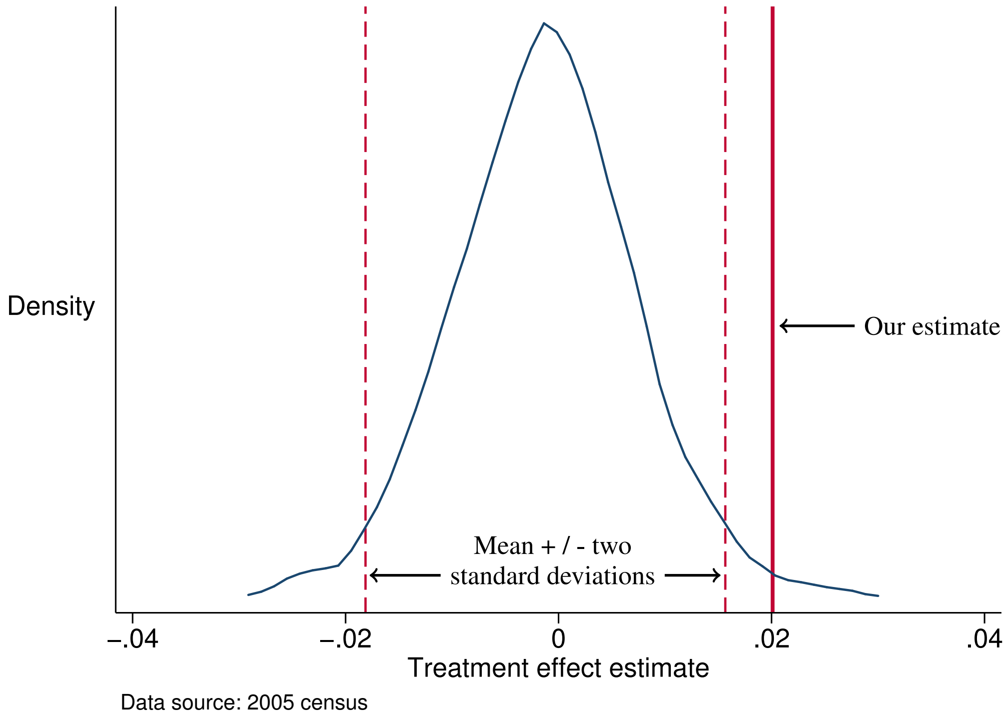

V.B. Mapping our estimate to the parameter of interest

In this section, we investigate how to map our treatment effect estimate to the parameter of greater interest - the labor market returns to the human capital accumulated in a year of schooling, isolated from the contribution of signaling. We noted in the introduction that our treatment effect estimate is a factor of four to ten lower than most other estimates of the returns to schooling from China and beyond (Ashenfelter, Harmon, and Oosterbeek, 1999; Card, 2001; Liu and Zhang, 2013). It is also much lower than the eight to 14 percent return per year we observe in the naive correlation between income and years of schooling in our data (results given in Table A.5). We claim that much of this gap between our estimate and the naive mincerian return is driven by our ability to shut down the signaling contribution to the returns to schooling by varying years of schooling while holding highest credential constant. This section provides several pieces of evidence suggesting that our treatment effect estimate is likely very close to the returns to schooling parameter that we are after.

First, we provide evidence that there was substantial learning in this extra year of primary school. We test for a difference between the treated and untreated in cognitive skills, as measured by a test administered to adult respondents in the CFPS survey. Figure VII plots the kernel density functions for treated and untreated individuals using the five year bandwidth sample. The

Figure VII: Kernel density of cognitive skills test score, by treatment status

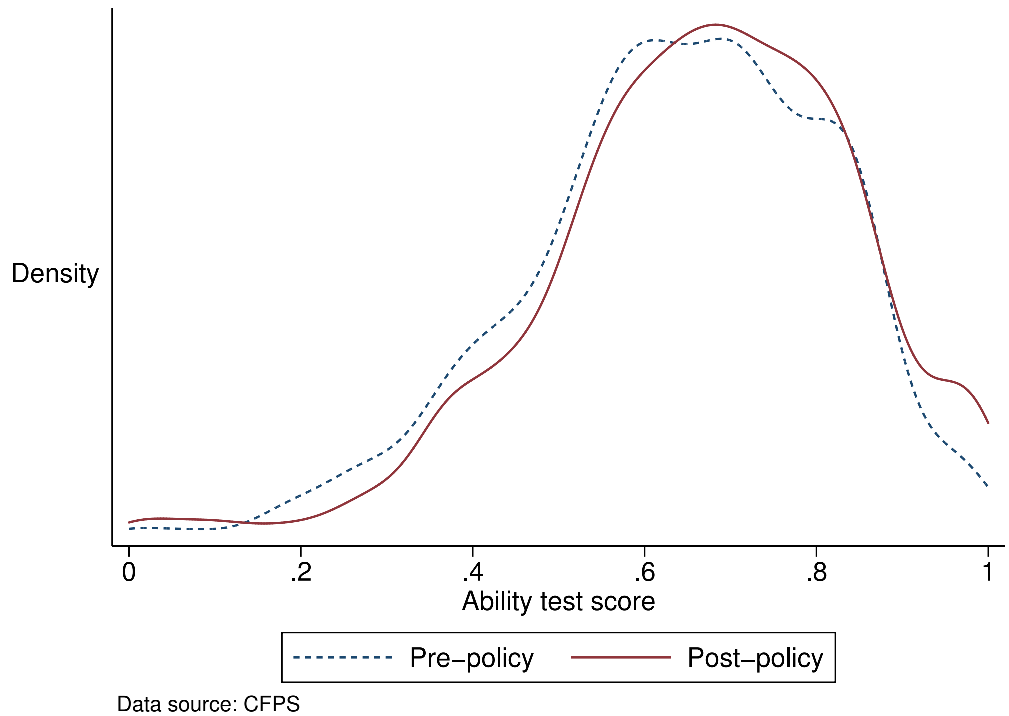

This plot shows the kernel density of cognitive skills test scores for observations from the CFPS data, plotted separately for those in affected and unaffected cohorts, in the five year bandwidth sample.

two distributions track each other quite closely, but there is a visible rightward shift in the treated distribution. A Kolmogorov-Smirnov test rejects the equality of the two distributions with a p-value of 0.003. The difference in distributions is most stark in the left tail, where previous work suggests we would be most likely to find it (Meghir and Palme, 2005; Dobbie and Fryer, 2013). The difference between the the left tails of the two distributions is 0.1 to 0.3 standard deviations, very similar to the estimated impact of a remedial primary education program in India (Banerjee et al., 2007). This result, alongside our estimate of the effect of the extra year on monthly income, adds to evidence that childhood interventions which may initially generate increases in test scores or cognitive skills often bring labor market returns decades after the initial intervention (Heckman, 2006; Chetty et al., 2011).

In addition, we observe larger estimates (i.e., larger income gains) among subgroups for whom we expect an extra year of learning to be particularly beneficial - those with less schooling, private sector workers, and women. If we assume diminishing returns to schooling, we would expect those with fewer years of schooling to enjoy a larger marginal benefit from a given improvement in human capital. Economic logic suggests private sector workers are likely to be paid more than government employees for a given increase in relative productivity. As discussed earlier, we also know from other work on China that the labor market returns to schooling are higher for women than men (Hannum, Zhang, and Wang, 2013).

The nature of the sixth year of primary school could also contribute to the gap between our estimates and the parameter of interest. While the fact that this additional year did not add new curriculum could potentially explain the small effect estimate, we note that the last year of both middle school and high school are primarily review for entrance exams to the next level of schooling, suggesting the extra year that we study is not so unusual in its content (Larmer, 2014). Our estimates of the returns to this sixth year may therefore be similar to the returns to these other years of schooling. In further support of this claim, we note that the only other well-identified estimate of the returns to schooling in China which attempts to disentangle the contributions of signaling and human capital is similar in magnitude. Li, Liu, and Zhang (2012) estimate that an additional year of schooling brings a 2.7 to 3.8% increase in earnings, using within-twin-pair differences in schooling to mitigate the potential for unobserved ability correlated with credential attainment (closely related to the contribution of signaling, as in Spence (1973) and Lange (2007)) to bias the results.

Second, we show there is little evidence of a gap between our treatment effect estimate, which estimates the effect of the policy on self-reported monthly income, and the true returns to schooling. There are four issues to report on here: mapping from reduced form to IV estimates, dealing with attenuation bias, the possibility that the general equilibrium effect of the policy is smaller than the parameter we are after, and estimating the range of potential downward bias stemming from the fact that the treated group has one less year of work experience than the untreated group due to the additional year spent in school.

Though it is common to "inflate" the reduced form treatment effect coefficient through dividing it by the proportion of individuals affected by the policy (the "first stage"), we argue that is inappropriate in our scenario for the following reasons. First, the proportion affected is unstable over bandwidth choice: the five year bandwidth estimate of the proportion of individuals induced to spend six years in primary school as a result of the policy is 0.55, while the 10 year estimate is almost 0.8. These estimates come from use of the CFPS data. Panel B of Figure V shows that the effect estimate of the policy on monthly income is stable across bandwidth choice, an estimate generated using the census data. This dispels concerns of downward bias because of either attenuation from misclassification of treatment status or a first stage substantially smaller than one. If such bias were to present in our estimates, we would expect the treatment effect estimate on income to increase with the number of years around the bandwidth we include in our sample, as this would increase the proportion of correctly identified observations. Given the stability we observe across the different bandwidth specifications for the coefficient estimate of the policy on income, we conclude the threat of this downward bias in our estimate is low.

Another potential explanation for this discrepancy is the difference between partial and general equilibrium effects. The policy we study affects a far larger proportion of the population than most studies which generate large effects, and it could be that the competition among many workers with the same ability gain drives down the labor market returns to the extra year. There are two reasons why we think this is unlikely to be the case. One, another study of an educational policy change which increases years of schooling for nearly half of the UK population finds much larger labor market returns (Oreopoulos, 2006). Two, we test for these general equilibrium effects and find no evidence of their existence[31].

Finally, we need to account for the fact that income increases with age and that some of this may be due to the returns to experience. As in all exercises attempting to estimate the returns to a year of schooling, we face the conundrum that an individual with one more year of schooling has one less year of experience in the labor market, and so our estimate may capture the returns to a year of schooling minus the returns to a year of experience. We calculate upper and lower bounds on this potential contribution, as in Manski (2013). The lower bound is 0, reflecting a world in which the returns to a year of age are entirely due to maturity. To calculate the upper bound, we assume that the difference between the average income of a 36 year old and a 35 year old with the same credential is entirely due to experience, and we calculate a weighted average of the per-year-of-age income premium we see in our data:

$$\sum_{age=t} \omega_t (\overline{Income_t} - \overline{Income_{t-1}}) \tag{3}$$

ω_t is the proportion of individuals in our estimation sample of age t. This generates an upper bound of 0.0168, suggesting that the upper bound of the range of estimates corrected for this bias would be 0.0371, still less than half of the lowest value in our naive mincerian per-year return estimates.

We can use these results and three assumptions to make a back-of-the-envelope calculation about the relative contribution of the signaling and human capital accumulation channels to the correlation between schooling and income. Our first assumption is that the returns to schooling flow only through two channels, the returns to skills gained and returns to the signal provided by the individual's highest educational credential. Our second assumption is that the final year of middle school and high school, each of which is also a review, yield the same return as our estimate of the return to the sixth year of primary school. Our final assumption is that the other years of schooling give twice the boost of the review years (i.e. 4.06%, slightly larger than the per-year estimate Li, Liu, and Zhang (2012) generate using twins for identification). Under these assumptions, we estimate that the signaling channel accounts for 57.2% of the returns to a middle school degree and 60.0% of the returns to a high school degree we observe in the cross-section.

VI. Interpretation and Discussion

This section of the paper has two parts. First, we perform a cost-benefit calculation of the policy's impacts. We finish the section with a discussion of the policy implications of our analysis.

[31] To run this test, we divide our sample into three groups based on when the policy was implemented: 1981-85, 1986-1990, and 1991-1995. The earlier the implementation, the more individuals exposed to the extra year of primary schooling, and so the closer the area is to the general equilibrium state of all workers being treated with this extra year. Though the earliest implementing group has a smaller treatment effect estimate than the later two groups (0.013, 0.033, and 0.031 for early, middle, and late implementers), consistent with the general equilibrium effect being smaller than the partial equilibrium effect, a Wald test fails to reject the equality of these three coefficients (p-value of the f-test: 0.32).

VI.A. Cost-benefit Analysis

We next use our results to generate an estimate of the net benefit of the program. We borrow our framework directly from Duflo (2001), focusing on the private gains and losses and ignoring the other potential benefits of increased wages (e.g. decreases in fertility and child morbidity). Though we are aware that this type of exercise involves a precarious amount of uncertainty (Manski, 2013), it is useful to gain some insight into the net effect of such a tremendous reallocation of resources.

As in Duflo's analysis, we choose our time frame to span from the first cohort in which some students leaving primary school are affected by the policy, 1981, to the end of 2050, and assume an annual discount rate of 5%. Our cost estimate has as its sole argument the lost year of wages[32], w_{it}, that affected students i forgo during the year t they spend in school instead of in the labor market

$$Cost = \sum_i w_{it} \tag{4}$$

For each cohort, we determine what proportion of individuals leave school with a primary, middle, high school or tertiary credential, and count the year lost as the last year they spent in school[33]. We then calculate the total value of the years lost for all students in each cohort from 1981 to the last cohort entering the labor force in 2050, using the same formula for the value of wages used in the benefit calculation below. Unlike Duflo, we do not incorporate a deadweight loss of taxation, as we assume there is no productive activity displaced by the policy other than the reallocation of affected individuals' time.

Our structure for estimating the benefit of this policy is also taken directly from Duflo's analysis. Specifically, we estimate the sum of income gains for all affected cohorts over the time frame we have chosen:

$$Benefit = \sum_t \sum_c \alpha GDP(t) S(c,t) P(c) \beta \tag{5}$$

Here α is the share of labor in GDP[34], $S(c,t)$ is the size of cohort c divided by the total working population in year t. P(c) is the proportion of cohort c affected by the policy, and β is our estimated effect of the policy on income. We sum the benefits earned by each cohort in the labor force in each year, assuming people work from when they leave school until age 65[35].

In Table V, we present four cost-benefit estimates for the period 1981-2050, varying two key assumptions about the nature of β. The first is whether to assign the average treatment effect to all individuals or to take into account the changing educational profile of the Chinese citizenry over this time and use the credential-specific treatment effects and data on the distribution of highest

[32] We assume that the extra year of primary school does not induce individuals to remain in the workforce for longer. In Appendix 4 we discuss other costs and our decision not to include them in this calculation.

[33] For example, an individual born in 1975 would start school in 1981. If her highest educational credential is a middle school degree and she was affected by the policy (and so she spent nine years in school), her "lost year" would be 1990.

[34] This labor share data comes from Karabarbounis and Neiman (2014). The rest of the data used in this section was downloaded from stats.gov.cn and the World Bank's World Development Indicators, projected forward using multi-year moving averages.

[35] This is a simplifying assumption. For those working in factories, the official retirement age is 60 for men and 50-55 women, but individuals often work well beyond these ages. In addition, the official age is slated to be changed in the next five years. Using the official retirement age would reduce the amount of years during which benefits accrue and thus reduce our estimate of the benefits of the policy.

Table V: Cost-benefit calculation: 1981-2050

	Assumption about heterogeneity in β_1 by highest credential held	
β_1 estimate used for rural areas	Using average treatment effect for all	Using credential-specific treatment effects
Using estimates from urban areas	-36,462	34,902
Using rural effect estimates from our data	-82,390	-212,220
Costs	178,497	-
Cost-benefit calculation for cohort leaving primary school in 2014 (urban β_1, average β_1)	-456.4	

*Estimates in millions of 2015 US Dollars.

educational credential in each cohort. The left column's figures use the average treatment effect estimates, while the figures in the right column use credential-specific estimates. The second is to decide how to estimate the effects for rural China. So far we have presented estimates for urban areas only, citing concerns about migration, the possibility for intensive or extensive reallocation of labor between agriculture and the non-farm rural labor market, and the great changes affecting rural China over this period. The policy, however, was implemented in both urban and rural areas and, until 2011, more than half of China's population was rural. To assess the likely fiscal impact of the policy on all of China, rural and urban, we present two sets of estimates. In the first row, we use the urban estimates on returns for both urban and rural residents. In the second row, we use the treatment effect estimates specific to urban and rural areas (the average treatment effect in the left column, then the credential-specific effects estimated off of only rural residents) weighted by the population share in a given year for the two groups.

Under three of the four scenarios, the policy's costs exceed its benefits by tens of billions of dollars. To provide an estimate of the current per-year cost of the policy, we compare the value of a lost year of productive work to the lifetime productivity benefits of the extra year of primary schooling for the cohort leaving primary school in 2014. We estimate this too to be a net loss of 456 million dollars. These figures underestimate the total costs, as we do not compute the value of parents' expenditures during that extra year on students' food, clothing, and other necessities. A conservative estimate of these would add hundreds of dollars to the per-person cost of the policy, which is tens of billions more in costs over the policy's lifetime.

VI.B. Policy Implications

We conclude this section with the policy implications of our research. Together, our results highlight the massive amount of resources that can be reallocated by a seemingly arbitrary policy choice: how many years should children spend in each level of school? To gain an idea of the generalizability of this finding, we analyzed the most recent DHS data from those 74 countries in which the number of years a respondent spends in school is collected. In 48 of these countries, we found evidence of bunching at multiple credential attainment years similar to that shown in Figure IV. We are less confident about the robustness of the DHS data to concerns about reporting error than we are for the Chinese data, but these patterns are consistent with our claim that government decisions on how long to make each level of schooling have massive resource implications[36].

Though we estimate the policy to be a net loss in China's case, it was redistributive: those with lower levels of schooling gained the most from it. This section's analysis reveals that bunching at credential attainment years is commonly found in developing countries. In such countries, a policy similar to the one we study which adds a year to the lowest level of schooling could be an effective way to ensure the less well-off gain skills valued in the labor market. The bleak cost-benefit estimates we generate in China's case are driven by the large increase in the proportion of individuals with secondary and tertiary credentials over time. In a country where the average number of years spent in school is lower and grows less quickly than it did in China between 1980 and 2015, the benefits of such a policy could well exceed the costs.

VII. Conclusion

In this paper, we exploited a massive policy change in China's educational system to study how household decisions on children's schooling respond to extending the length of primary school by one year. We observe extensive bunching at credential attainment years both before and after the policy. The vast majority of affected individuals spent an additional year in school, sacrificing an entire year of earnings, in order to earn the credential they would have attained in the absence of the policy. These results allowed us to generate a new parameter estimate of the returns to a year of schooling, holding highest educational credential constant. We found a small but precisely estimated two percent increase in monthly income which is higher for individuals with less schooling, those whom the policy set out to assist.

We then showed that our estimate of the effect of the extra year of primary school on income is quite close to a parameter that has long been of interest to labor economists: the labor market returns to the human capital accumulated in schooling independent of the contribution from the signaling mechanism. Finally, we estimated the costs and benefits of the nearly one trillion person-hours this policy has reallocated from the labor market to the pursuit of schooling, finding that in most scenarios the policy generates a net loss of tens of billions of dollars.

[36]Countries in which we observe bunching at credential attainment years had more than twice the per-capita GDP than those without bunching. If productive characteristics for non-agricultural occupations are less observable than those for agricultural work (under the assumption it is more difficult to observe brains than brawn), the greater signaling among richer countries we observe is consistent with a larger contribution of signaling to schooling returns in economies with more developed non-agricultural labor markets.

In our policy analysis, we argued that our paper shines light on an important research lacuna: government policy which sets the length of each level of schooling is under-researched given the amount of resources at stake. In addition to the value of the time these policies reallocate, over three trillion US dollars are spent worldwide on education by national governments each year (World Bank, 2012). Essentially every government on the planet must decide how long nationally-sanctioned school will last and how much to spend on it. Providing more informed guidance on how to make these decisions is an important next step for education research in economics.

REFERENCES

Almond, Douglas, Hongbin Li, and Shuang Zhang. 2013. "Land reform and sex selection in China." *NBER Working Paper* .

Angrist, Joshua and Alan Krueger. 1991. "Does Compulsory School Attendance Affect Schooling and Earnings?" *Quarterly Journal of Economics* 106 (4):979–1014.

Ashenfelter, Orley, Colm Harmon, and Hessel Oosterbeek. 1999. "A review of estimates of the schooling/earnings relationship, with tests for publication bias." *Labour Economics* 6 (4):453–470.

Banerjee, Abhijit and Esther Duflo. 2011. *Poor economics: A Radical Rethinking of the Way to Fight Global Poverty*. Public Affairs.

Banerjee, Abhijit V., Shawn Cole, Esther Duflo, and Leigh Linden. 2007. "Remedying education: Evidence from two randomized experiments in India." *Quarterly Journal of Economics* 122 (3):1235–1264.

Bedard, Kelly. 2001. "Human capital versus signaling models: university access and high school dropouts." *Journal of Political Economy* 109 (4):749–775.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How much should we trust differences-in-differences estimates?" *Quarterly Journal of Economics* 119 (1):249–275.

Black, Sandra E, Paul J Devereux, and Kjell G Salvanes. 2005. "Why the Apple Doesn't Fall Far: Understanding Intergenerational Transmission of Human Capital." *American Economic Review* 95 (1):437–449.

Butrymowicz, Sarah. 2012. "In China, private colleges, universities multiply to meet higher-education demand." *Washington Post* .

Cai, Fang, Albert Park, and Yaohui Zhao. 2008. "The Chinese labor market in the reform era." In *China's Great Economic Transformation*. Cambridge University Press.

Card, David. 2001. "Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems." *Econometrica* 69 (5):1127–1160.

Chetty, Raj, John N Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, and Danny Yagan. 2011. "How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project Star." *The Quarterly Journal of Economics* 126 (4):1593–1660.

Chyi, Hau and Bo Zhou. 2014. "The effects of tuition reforms on school enrollment in rural China." *Economics of Education Review* 38:104–123.

Clark, Damon and Paco Martorell. 2014. "The signaling value of a high school diploma." *Journal of Political Economy* 122 (2):282–318.

Clark, Damon and Heather Royer. 2013. “The Effect of Education on Adult Mortality and Health: Evidence from Britain.” *American Economic Review* 103 (6):2087–2120.

Connelly, Rachel and Zhenzhen Zheng. 2003. “Determinants of school enrollment and completion of 10 to 18 year olds in China.” *Economics of Education Review* 22 (4):379–388.

Dobbie, Will and Roland G Fryer. 2013. “Getting Beneath the Veil of Effective Schools: Evidence From New York City.” *American Economic Journal: Applied Economics* 5 (4):28–60.

Duflo, Esther. 2001. “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment.” *American Economic Review* 91 (4):795–813.

Fang, Hai, Karen N. Eggleston, John A. Rizzo, Scott Rozelle, and Richard J. Zeckhauser. 2012. “The Returns to Education in China: Evidence from the 1986 Compulsory Education Law.” *NBER Working Paper* (18189).

Fredriksson, Peter, Björn Öckert, and Hessel Oosterbeek. 2013. “Long-Term Effects of Class Size.” *Quarterly Journal of Economics* 128 (1):249–285.

Fukunaga, Keinosuke and Larry Hostetler. 1975. “The estimation of the gradient of a density function, with applications in pattern recognition.” *IEEE Transactions on Information Theory* 21 (1):32–40.

Gelman, Andrew and Guido Imbens. 2014. “Why high-order polynomials should not be used in regression discontinuity designs.” *NBER Working Paper* .

Grenet, Julien. 2013. “Is Extending Compulsory Schooling Alone Enough to Raise Earnings? Evidence from French and British Compulsory Schooling Laws.” *Scandinavian Journal of Economics* 115 (1):176–210.

Hannum, Emily. 1999. “Political change and the urban-rural gap in basic education in China, 1949-1990.” *Comparative Education Review* 43 (2):193–211.

Hannum, Emily, Yuping Zhang, and Meiyan Wang. 2013. “Why Are Returns to Education Higher for Women than for Men in Urban China?” *The China Quarterly* 215:616–640.

Hannum, Emily C, Jere Behrman, Meiyan Wang, and Jihong Liu. 2008. “Education in the reform era.” In *China’s great economic transformation*. Cambridge University Press.

Heckman, James J. 2006. “Skill formation and the economics of investing in disadvantaged children.” *Science* 312 (5782):1900–1902.

Imbens, Guido W and Thomas Lemieux. 2008. “Regression discontinuity designs: A guide to practice.” *Journal of Econometrics* 142 (2):615–635.

Karabarbounis, Loukas and Brent Neiman. 2014. “The Global Decline of the Labor Share.” *The Quarterly Journal of Economics* 129 (1):61–103.

Kipnis, Andrew B. 2011. *Governing educational desire: Culture, politics, and schooling in China*. University of Chicago Press.

Lang, Kevin and David Kropp. 1986. "Human capital versus sorting: the effects of compulsory attendance laws." *Quarterly Journal of Economics* 101 (3):609–624.

Lange, Fabian. 2007. "The speed of employer learning." *Journal of Labor Economics* 25 (1):1–35.

Larmer, Brook. 2014. "Inside a Chinese Test-Prep Factory." *New York Times* .

Lee, David S and David Card. 2008. "Regression discontinuity inference with specification error." *Journal of Econometrics* 142 (2):655–674.

Lee, David S and Thomas Lemieux. 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature* 48 (2):281–355.

Li, Haizheng. 2003. "Economic transition and returns to education in China." *Economics of Education Review* 22 (3):317–328.

Li, Hongbin, Lei Li, Binzhen Wu, and Yanyan Xiong. 2012. "The end of cheap Chinese labor." *Journal of Economic Perspectives* 26 (4):57–74.

Li, Hongbin, Pak Wai Liu, and Junsen Zhang. 2012. "Estimating returns to education using twins in urban China." *Journal of Development Economics* 97 (2):494–504.

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.

Liu, Elaine and Shu Zhang. 2013. "A Meta-Analysis Of The Estimates Of Returns To Schooling In China." *Working Paper* .

Liu, Yingjie. 1993. *Book of Major Educational Events in China 1949-1990 (in Chinese)*. Zhejiang Education Publishing House.

Loyalka, Prashant, Chengfang Liu, Yingquan Song, Hongmei Yi, Xiaoting Huang, Jianguo Wei, Linxiu Zhang, Yaojiang Shi, James Chu, and Scott Rozelle. 2013. "Can information and counseling help students from poor rural areas go to high school? Evidence from China." *Journal of Comparative Economics* 41 (4):1012–1025.

Lucas, Adrienne M and Isaac M Mbiti. 2012. "Access, sorting, and achievement: the short-run effects of free primary education in Kenya." *American Economic Journal: Applied Economics* 4 (4):226–253.

Lv, Ping and Yu Xie. 2012. "Sampling Design of the Chinese Family Panel Studies." Working paper, Institute of Social Sciences Surveys, Peking University. URL `http://www.uchicago.cn/wp-content/uploads/2011/05/Ping-Lv.pdf`.

Manski, Charles F. 2013. *Public policy in an uncertain world: analysis and decisions*. Harvard University Press.

McCrary, Justin. 2008. "Manipulation of the running variable in the regression discontinuity design: A density test." *Journal of Econometrics* 142 (2):698–714.

Meghir, Costas and Mårten Palme. 2005. "Educational reform, ability, and family background." *American Economic Review* 95 (1):414–424.

Ministry of Education of the People's Republic of China, Beijing. 2013. "China Education Statistical Yearbooks." URL `http://www.moe.gov.cn/publicfiles/business/htmlfiles/moe/s7255/201303/149854.html`.

Mo, Di, Linxiu Zhang, Hongmei Yi, Renfu Luo, Scott Rozelle, and Carl Brinton. 2013. "School Dropouts and Conditional Cash Transfers: Evidence from a Randomised Controlled Trial in Rural China's Junior High Schools." *Journal of Development Studies* 49 (2):190–207.

Mok, Ka Ho. 2000. "Marketizing higher education in post-Mao China." *International Journal of Educational Development* 20 (2):109–126.

Munshi, Kaivan and Mark Rosenzweig. 2013. "Networks, Commitment, and Competence: Caste in Indian Local Politics." *NBER Working Paper* .

National Institute, of Education Sciences. 1984. *Chronicle of Education Events in China (in Chinese)*. Educational Science Publishing House.

Oreopoulos, Philip. 2006. "Estimating average and local average treatment effects of education when compulsory schooling laws really matter." *American Economic Review* 96 (1):152–175.

Shi, Xinzheng. 2012. "Does an intra-household flypaper effect exist? Evidence from the educational fee reduction reform in rural China." *Journal of Development Economics* 99 (2):459–473.

Spence, Michael. 1973. "Job market signaling." *Quarterly Journal of Economics* 87 (3):355–374.

Stephens, Melvin and Dou-Yan Yang. 2014. "Compulsory education and the benefits of schooling." *American Economic Review* 104 (6):1777–1792.

Tyler, John H, Richard J Murnane, and John B Willett. 2000. "Estimating the labor market signaling value of the GED." *Quarterly Journal of Economics* 115 (2):431–468.

UNICEF. 1978. *Statistics on children in UNICEF assisted countries*. Unicef.

Vogel, Ezra F. 2011. *Deng Xiaoping and the transformation of China*. Belknap Press of Harvard University Press.

Weiss, Andrew. 1995. "Human capital vs. signalling explanations of wages." *The Journal of Economic Perspectives* :133–154.

World Bank, Publications. 2012. *World Development Indicators 2012*. World Bank Publications.

Zhang, Junsen, Yaohui Zhao, Albert Park, and Xiaoqing Song. 2005. "Economic returns to schooling in urban China, 1988 to 2001." *Journal of Comparative Economics* 33 (4):730–752.

APPENDIX

APPENDIX 1 - APPENDIX FIGURES

Figure A.1: National data on proportion of students in six year primary education

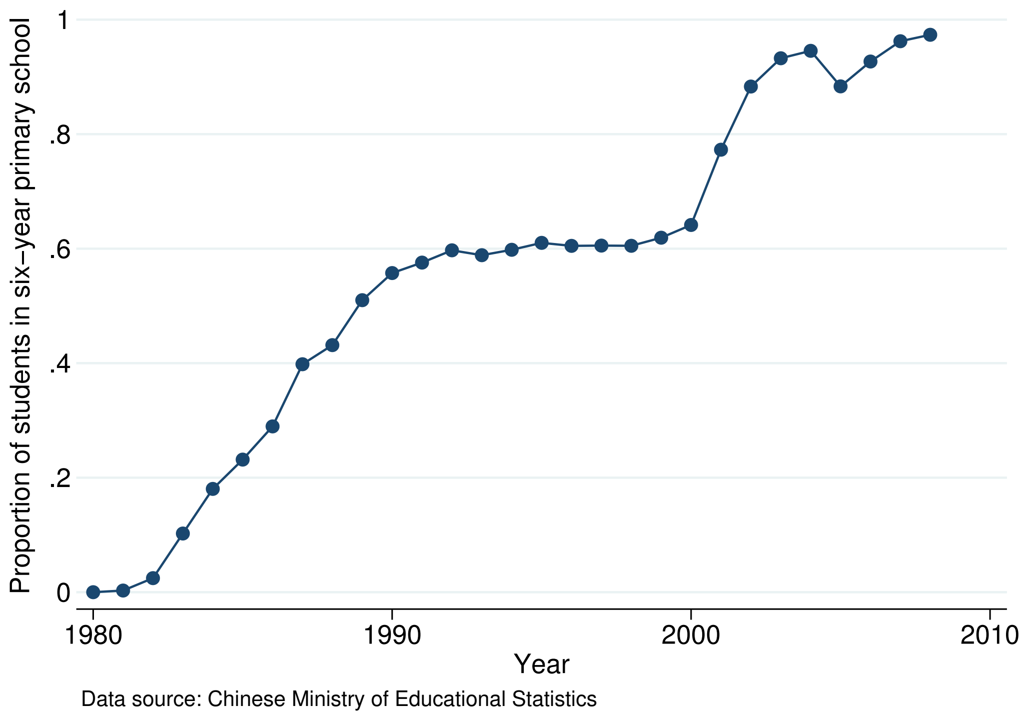

Figure A.2: Timing of implementation by prefecture

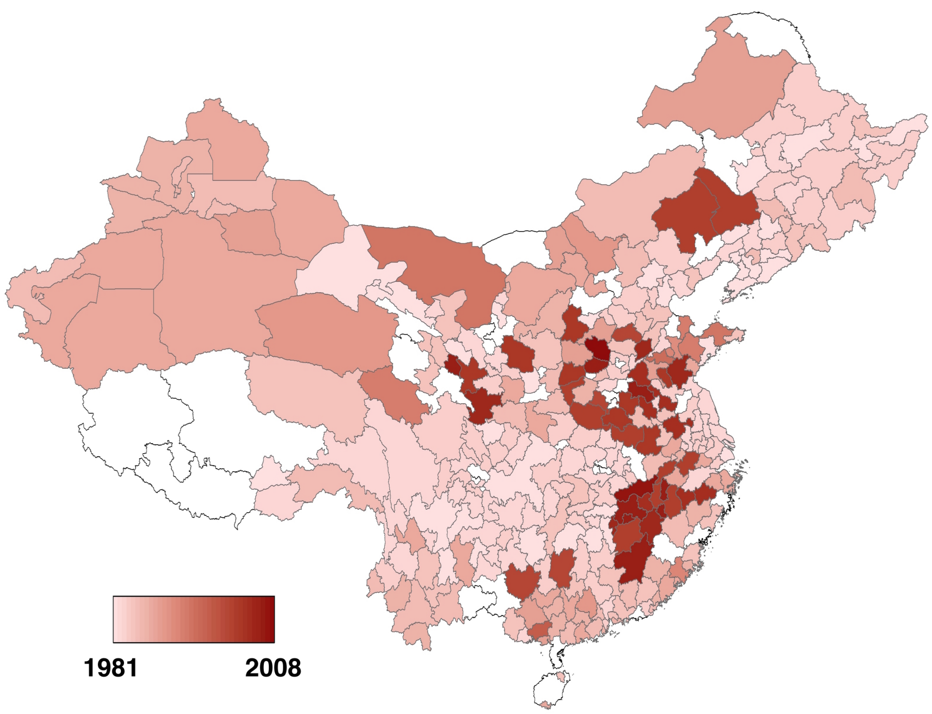

Note: In some cases, prefecture boundaries have changed since the archive was published. In these cases, we take the prefecture level-average of the treatment year in all previous prefecture capitals falling within the new prefecture. In prefectures that are colored white, we are unable to assign a unique prefecture-level treatment year for reasons discussed in Section III.

Figure A.3: Example of a page from a county gazetteer

续表

1979	440	85 846	3 721
1980	355	84 274	3 565
1981	362	83 952	3 604
1982	386	81 925	3 509
1983	353	78 492	3 328
1984	361	80 059	3 373

2．学制、课程设置及教法

1976年，全县小学实行五年制。1984年，按照省教育厅要求，将小学五年制改为六年制，并在两年内完成。按照教育部门的要求，境内小学采取农村以乡，城镇以校为单位自行命题进行考试、择优的办法，使五年级应届毕业生，三分之二毕业升入初中。1985年应届毕业生三分之一升入初中，留校三分之二升入六年。

1976年秋，学校所开课程有语文、数学、自然、地理、历史、音乐、美术、体育、劳动，使用全国统编教材。1982年增设思想品德课。1984年，仍延用原有的课程设置及教材。

1976年后，各小学校逐步恢复正常教学秩序。1979年，县革委会发出了55号文件，摘掉了小学附设的初中班。恢复了中心校管理体制。1981年，学校工作的立足点转为以教学为中心，注重提高教学质量，加强教学研究和改革，强调“双基”（ 基础知识、基本技能）教学。要求教师立足课堂，提高课堂教学效果，着眼点放在培养学生能力、开发学生智力上。

1984年，各校进行教学改革试验，坚持以学生为主体，以教师为主导，全面进行教学工作整体改革。按照教育局要求，各校成立了课外活动小组，开展了“第二课堂”活动。学校在加强基础知识教学同时，还开展了小学语文、数学知识竞赛。是年，全县小学毕业生双科合格率达90%。

第二节　教育体制改革后的小学教育

一、学校发展概况

1985年，全县有小学校335所，其中教育部门和集体办320所，其他部门办15所，设立下伸点65个（教育部门和集体办58个，其他部门办7个）。在校学生80 464人（女学生38 583人），其中教育部门和集体办76 135人，其他部门办4 329人。教学班2 431个，其中教育部门办和集体办2 294个（复式班105个），其他部门办137个（复式班2个）。二部制学校7所，有教学班42个。教职工3 653人，其中教育部门办1 848人，集体办1 605人，其他部门办200人。

1987年，林海、三合乡划归白城市（现洮北区）后，境内有学校318所（小学中心校28所），教学点58个。在校学生74 992人，教学班2 334个（复式班100个，二部制班56个），教职工3 517人。

1989年，城内小学建成2座教学楼，建筑面积6 460平方米。其中，实验小学教学楼建筑面积3 260平方米；第四小学教学楼建筑面积3 200平方米。

Figure A.4: Example of mean shift algorithm identifying year of policy implementation

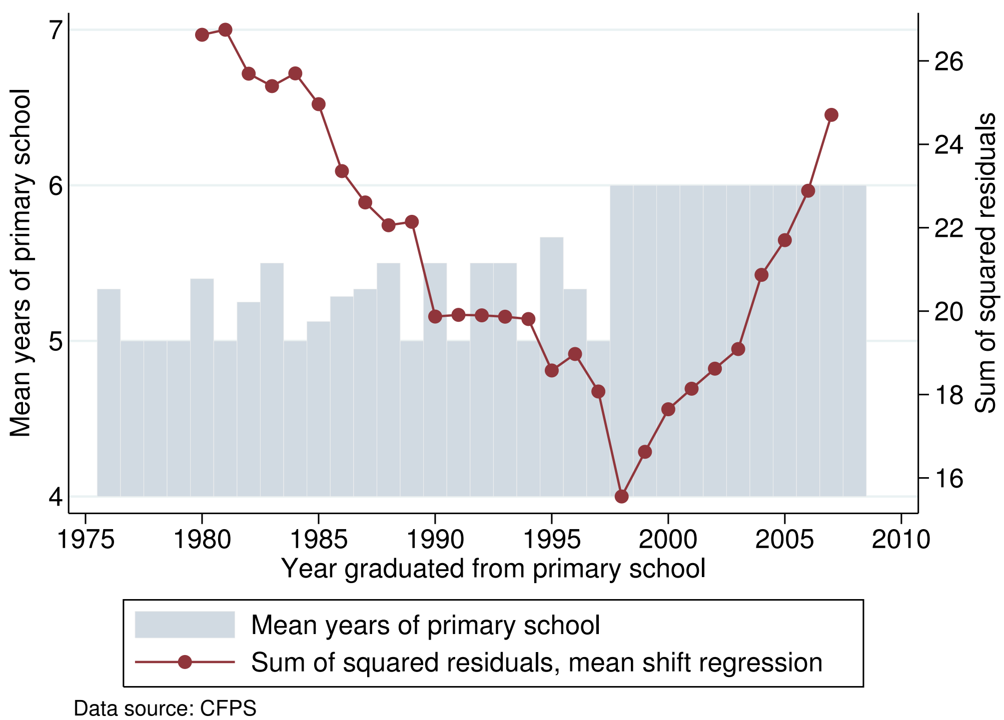

Figure A.5: Event study plots

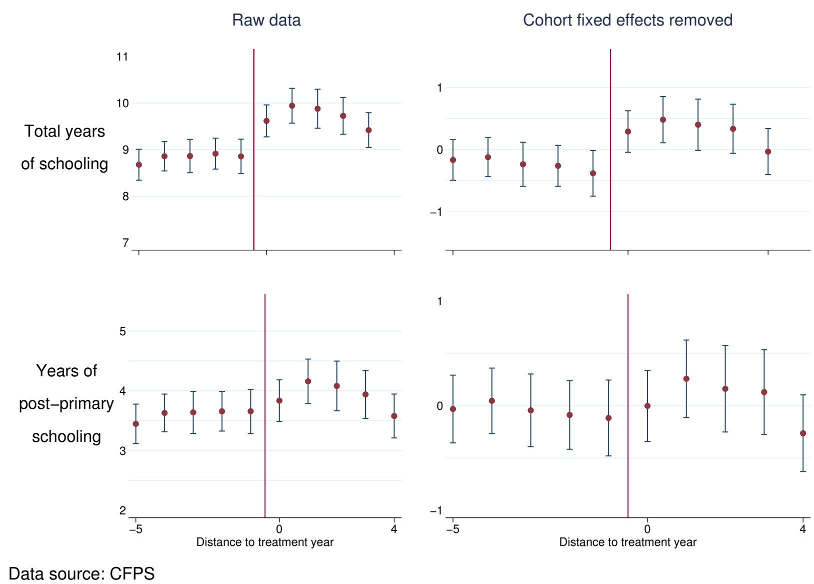

Panel A: Total years of schooling and years of post-primary schooling

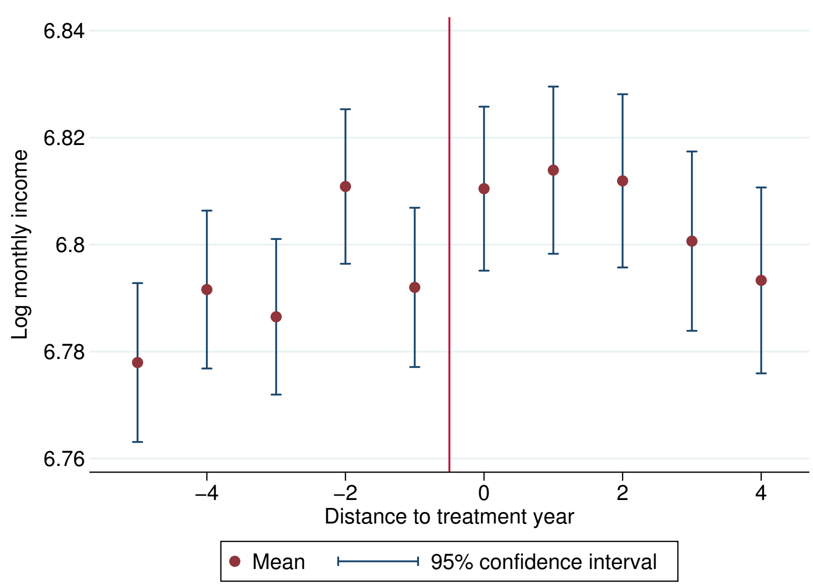

Panel B: Log monthly income

Figure A.6: Evidence of students' ability to adjust on lower and upper margins

Panel A: Ability to adjust on lower margin

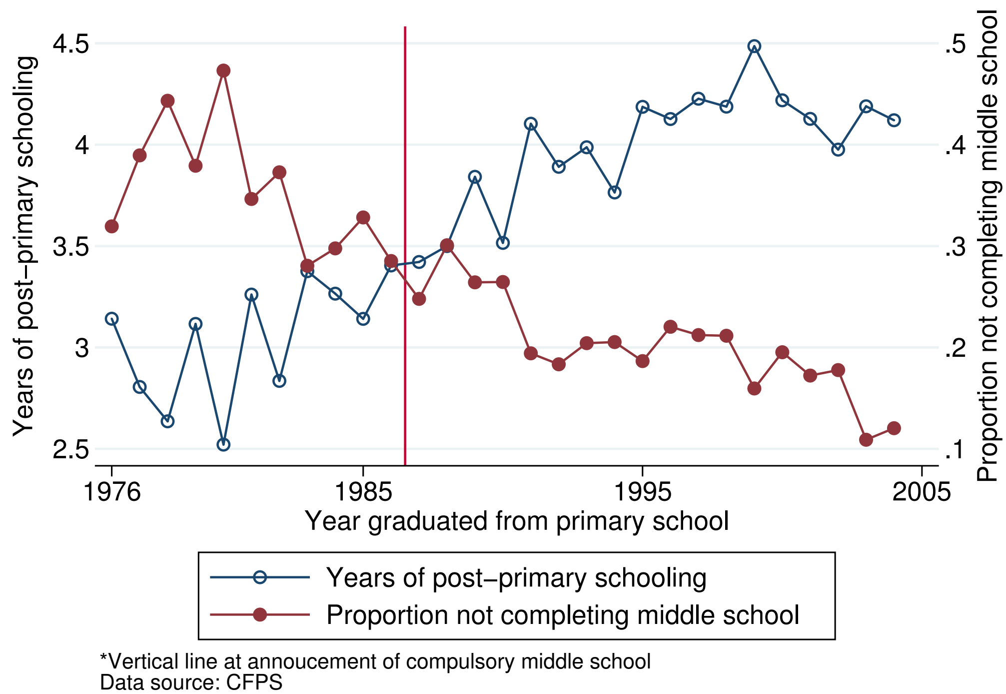

Panel B: Ability to adjust on upper margin

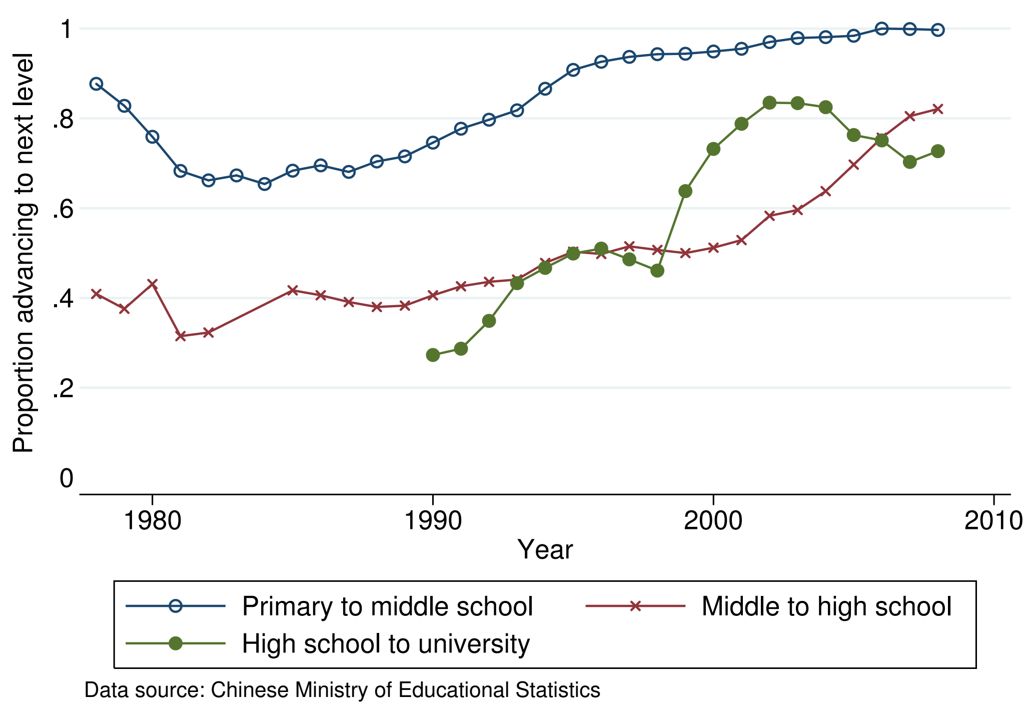

Figure A.7: National statistics and province-level CDFs of policy implementation

First 12 provinces

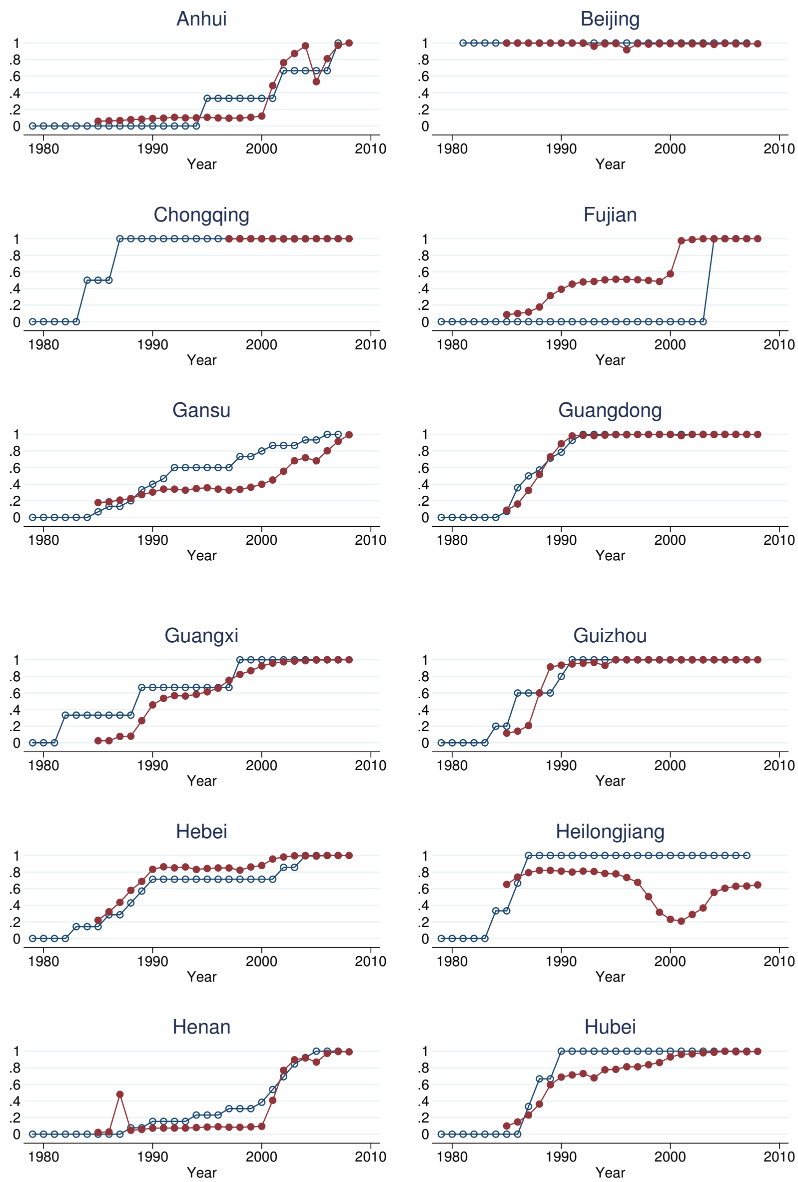

These plots show the proportion of individuals enrolled in six-year primary schools in the 24 non-Shanghai CFPS provinces. The sequence of hollow circles represent the proportion of treated counties in a given province by year according to county-specific implementation years identified in the CFPS data. The sequence of solid circles are from publicly available, Chinese Ministry of Education statistics on the number of schoolchildren benefiting from six year education in a given province by year. Measurement error and policy experimentation cause these to briefly trend downward in a few cases.

Figure A.8: National statistics and province-level CDFs of policy implementation

Second 12 provinces

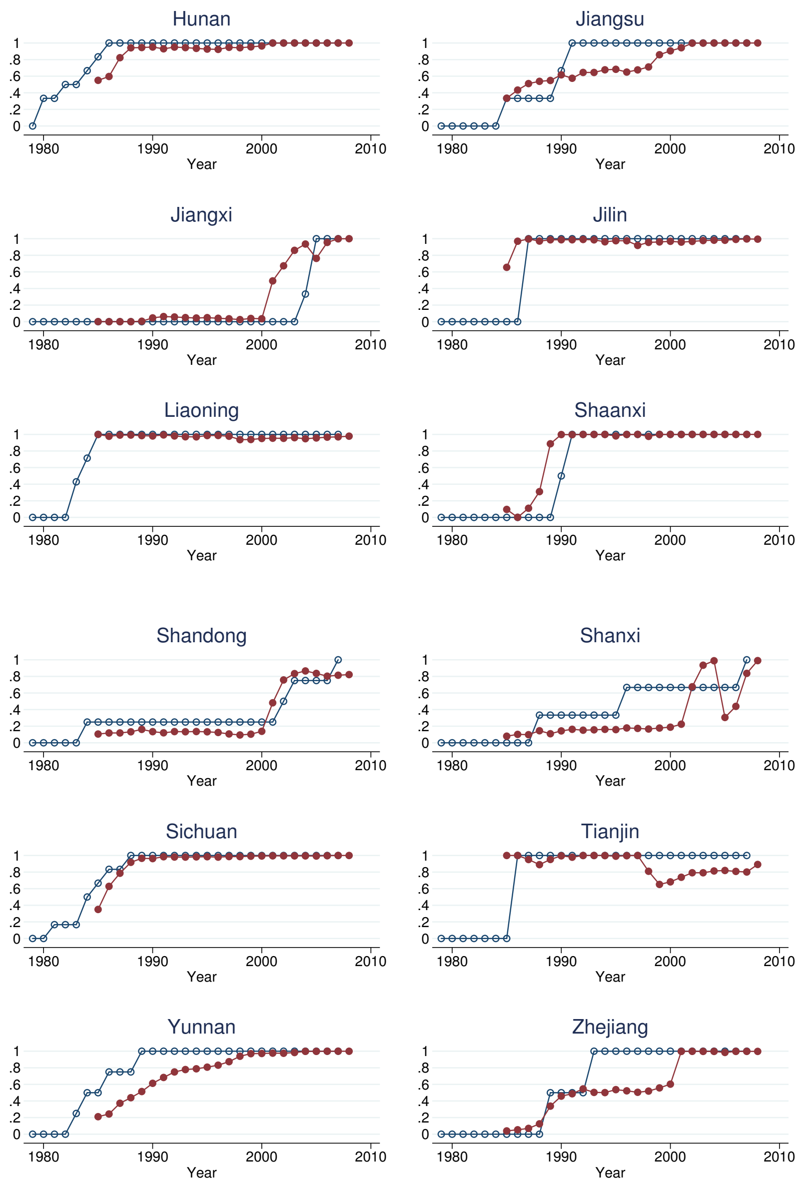

These plots show the proportion of individuals enrolled in six-year primary schools in the 24 non-Shanghai CFPS provinces. The sequence of hollow circles represent the proportion of treated counties in a given province by year according to county-specific implementation years identified in the CFPS data. The sequence of solid circles are from publicly available, Chinese Ministry of Education statistics on the number of schoolchildren benefiting from six year education in a given province by year. Measurement error and policy experimentation cause these to briefly trend downward in a few cases.

Figure A.9: Comparing archival and algorithmically identified treatment years

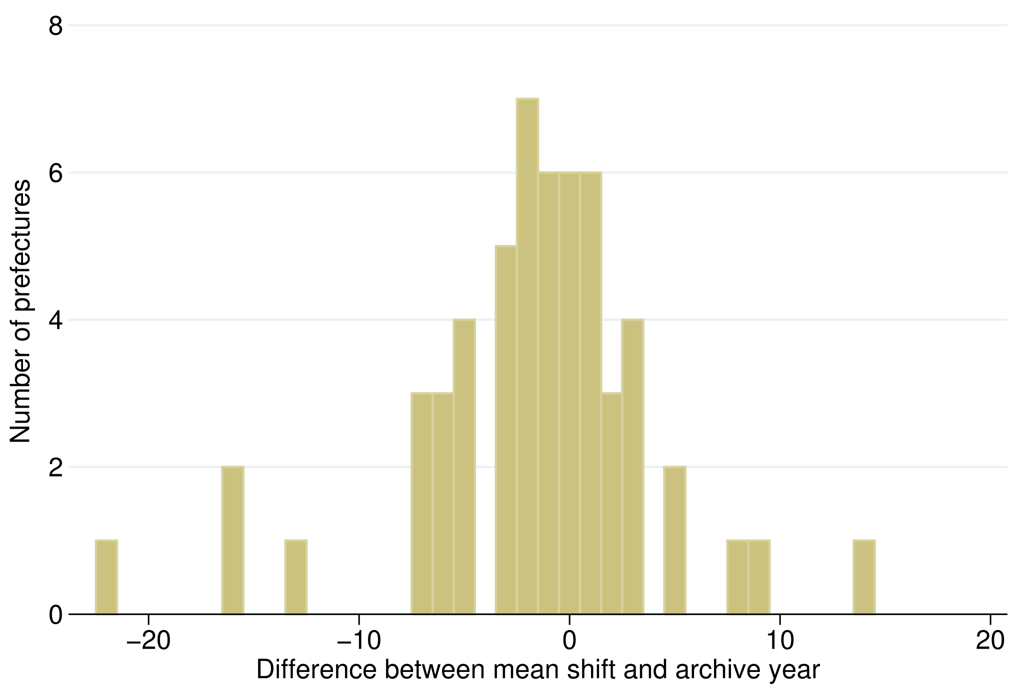

Data source: CLDS

Table A.1: Data sources

Data source	Year(s) collected	Sample size	Relevant contents	Policy year identification method	Geographic specificity
County educational gazetteers	Various	345*	Date policy implemented, implementation details	N/A	County**
Mini-census	2005	2.59m	Schooling: highest credential, dropout Labor market: income, employment	Archival records	Prefecture**
China Family Panel Studies (CFPS)	2010	33,600	Schooling: years spent in each level of schooling Labor market: income, employment	Algorithm	County
China Labor-Dynamics Survey (CLDS)	2012	16,253	Schooling: years spent in each level of schooling	Algorithm force and archival records	Prefecture

*Number of prefectures for which we collect gazetteer data
**Note: There are multiple counties in each prefecture

Table A.2: Anecdotes on implementation of six year primary education from across China

Locality	Strategy of Changing 5-Year to 6-Year Primary Education	Source
Beijing	In 1969, the length of primary education was shortened to five years. Starting from September 1st, 1980, the length of primary education was prolonged to six years, with the fifth-grade students continuing to be in the old system and students of other grades entering into the new system.	Beijing General Education Chronicle (Part I)
Xinle County, Hebei Province	In 1967, the length of primary education was shortened to five years. Starting from August 1985, first-grade students and one half of second-grade students entered into the new six year system, while the rest of students remained in the old system.	Xinle Educational Chronicle
Nanjing City, Jiangsu Province	In 1969, the length of primary education was shortened to five years. Starting from 1982, the length of primary education was prolonged to six years in urban districts and children started primary education at the age of six. By contrast, the length of primary education remained to be five years until 1999 in five other counties (Jiangning, Jiangpu, Liuhe, Gaochun, and Lishui).	Nanjing Educational Chronicle
Wuyi County, Zhejiang Province	In 1984, first-grade students entered into the new six year system in half of primary schools, with the other half following the next year. In 1987, however, all primary schools were required to resume the five year system. In September 1999, all primary schools were restored to the new system of six years. The final cohort under the old system graduated in June 2004.	Wuyi Educational Chronicle
Dongying District, Shandong Province	In 1997, the compulsory education system changed from the 5-3 (years of primary school-years of middle school) to the 5-4 system. In 2003, first-grade students entered into the new 6-3 system.	Dongying District Chronicle (1998-2005)
Xishui County, Hubei Province	In 1986, the first-grade students entered into the new 6-3 system in the primary schools located in county seats, while other primary schools remained in the old 5-3 system. In 1987, the first-grade students in the primary schools located in the township entered into the new 6-3 system. In 1991, first-grade students in the remaining primary schools entered into the new 6-3 system.	Xishui Educational Chronicle (1986-2006)

Table A.3: Summary statistics

	Data set			
	Census		CFPS	
Proportion	Rural	Urban	Rural	Urban
Demographics				
Female	0.518	0.517	0.458	0.494
Ethnic minority	0.172	0.070	0.081	0.050
Highest credential				
Middle school	0.575	0.337	0.489	0.417
High school	0.060	0.319	0.157	0.231
Tertiary	0.007	0.302	0.029	0.141
Labor market characteristics				
Currently working	0.909	0.780	0.597	0.693
Private sector	0.181	0.317	0.151	0.320
Entrepreneur	0.074	0.120	0.383	0.195
Observations	157,308	86,240	1,293	1,387

This table presents descriptive statistics of demographic, educational, and labor market characteristics for our two main datasets using the five-year bandwidth estimation sample. This is the same sample used to generate the regression estimates presented in the paper's main tables.

Table A.4: Regression of prefecture characteristics on timing of treatment

Variable	Coefficient	Standard error
Proportion male	-1.281	6.558
Proportion minority	0.023	0.751
Proportion urban	0.020	1.352
Proportion with government job	1.700	3.515
Mean log income	-0.322	0.657
Mean education level (1-5 scale)	0.122	0.829
Central region	-1.340***	0.472
Northeast region	-1.940***	0.503
Western region	-0.323	0.503
Treatment year summary statistics	1986.1	2.229

Data: census. Robust standard errors in parentheses.
N=240 prefectures. * p<0.10, ** p<0.05, *** p<0.01

Table A.5: Putting our estimates in context

Credential	(1) Credential premium	(2) Per-year return	(3) Percent employed	(4) Percent of estimation sample
Primary school	-	-	62.3	3.0
Middle school	23.7	7.93	66.9	28.9
High school	25.3	8.42	77.7	31.8
Technical college	25.7	8.58	92.1	22.8
University	55.7	13.9	95.8	13.4
Our estimate of the return to a sixth year of primary school	-	1.9-2.0	-	-
Li et al. (2012) twins estimate of the returns to year of schooling	-	2.7-3.8	-	-

N = 85,048. Data: census. Sample: 5-year bandwidth, employed urban residents. Column (1) gives the naive estimate of the total income bonus from gaining a given credential. The degree premia are: middle school relative to primary earners, high school relative to middle, and technical college and university both relative to high school. Column (2) presents the per-year premium to each credential, calculated by dividing the value in Column (1) by the number of years it takes to earn that credential.

Appendix 3 - Details on Our Use of the Mean Shift Model

As we use an algorithm to determine treatment year status for observations in the CFPS, we introduce a few safeguards to ensure that we are not erroneously choosing a year when there is no change in treatment status or not enough evidence to determine the year. Specifically, we exclude counties in which:

1. The cohort-level median number of years of primary school in a given county is never less than 5.5 before the identified treatment year or is never more than 5.5 after;

2. The mean shift model predicts a treatment year after which individuals have less primary schooling, that is, $\gamma_1 < 0$ in equation 1;

3. There are more than four cohorts for which we have no observations within five years before or after the identified treatment year (that is, data for at least five of the ten cohorts in this county's restricted bandwidth sample are missing).

This exercise excludes 32 of the 144 non-Shanghai counties, or about 78% of the sample. All of the results we show are qualitatively similar with the entire 144 county sample and are available by request.

As a final check of accuracy, we visually compare the treatment year generated by the mean shift model to each county's histogram. In four cases, our inspection of the histograms suggested a year more than one year different than the mean shift model's choice and we use this visually identified year in our main analyses, controlling for those counties whose year was chosen visually as opposed to algorithmically.

To check that these algorithmically identified years are credible, we conduct two exercises. First, we compare the mean shift years for all counties within each province to national statistics on provincial-level implementation, shown in Figures A.7 and A.8. The two series track closely and their correlation is 0.7759. Second, we apply both the mean shift and archival match policy year identification methods to the China Labor-force Dynamics Survey data. Figure A.9 shows the distribution of the difference between archival and algorithmically identified treatment years. These values are normally distributed around 0, and in 67% of prefectures the years identified by the two methods are within 3 years of each other. This amount of discrepancy is not surprising, as in the gazetteers we see that counties occasionally implemented at different times within a prefecture. Additionally, the number of observations in cohort-by-prefecture bins in the Labor-force Dynamics dataset is small, which implies imprecision (relative to using the CFPS) in our ability to algorithmically determine the implementation year.

Alternative Explanations for Bunching at Credential Attainment Years

In our theoretical framework, we abstract from several phenomena which could generate bunching at credential attainment years and inelasticity of demand for credentials with respect to credential length in the absence of credentialism. The most prominent of these potential confounders are compulsory schooling laws, price differentials between levels of schooling, supply constraints, high human capital accumulation in credential attainment years, the lottery value of sitting a high school or college entrance exam, the cultural value of credential attainment, and reporting bias. In this section we deal with each in turn, providing evidence which suggests that few of these issues are of concern and none can account for more than a fraction of the bunching we set out to explain.

China's 1986 compulsory education law stipulated that all Chinese students had to complete primary and middle school. While the policy could conceivably have generated much of the bunching at the middle school credential attainment year we observe, we estimate that it is unlikely to have done so. Recent empirical work shows the law was gradually implemented and remains porously enforced, with many places not implementing it until the mid-2000's (Fang et al., 2012). Using province-specific policy implementation years from that study, we estimate the policy's rollout has no effect on whether individuals attain a middle school credential or on the number of years they spend in primary school[37]. We believe this null result is driven by porous enforcement of the law in rural areas and the fact that, in urban areas, high educational attainment levels (as seen in our data) make the law non-binding for the vast majority of affected individuals. Further supporting the claim of porous enforcement in rural areas, recent work has shown that as recently as 2010 the proportion of students not completing middle school was over 20% in some areas of rural China (Mo et al., 2013). In Panel A of Figure A.6, we show that the cohort-specific proportion of students who get less than a full middle school education (three years of post-primary schooling) is stable at around 20% throughout most of the period we study.

The increase in school fees students face as they advance to each higher level of schooling is also unlikely to explain much of the bunching we observe. Though there is evidence that these school fee jumps pose binding constraints for China's rural poor (Liu et al., 2009), it appears that cost differentials do not determine schooling levels for the majority of the Chinese population. Reductions in school fees have little impact on overall enrollment, even in rural China (Shi, 2012; Chyi and Zhou, 2014), and there is an abundance of qualitative accounts documenting that the seats in Chinese universities are oversubscribed (Kipnis, 2011; Butrymowicz, 2012; Larmer, 2014). We test for a relationship between the cost of schooling and enrollment, using our main estimating equation and the introduction of school fees at the tertiary level in 1995 as plausibly exogenous variation in the cost gap between levels of schooling (Mok, 2000). We find no evidence that the increase in cost at this time is associated with a change in the proportion of individuals advancing to tertiary schooling.

Supply constraints are another possible driver of bunching at credential attainment years. While this explanation surely accounts for some of the bunching we observe, the importance of demand exceeding supply in generating bunching at credential attainment years is decreasing over the duration of our study. As China was continuously expanding its supply of schooling over the study period, the fertility reduction policies of the late 1970's and early 1980's reduced the number of individuals in a cohort competing for a given number of spots in school. National-level data plotted in Panel B of Figure A.6 shows that supply constraints are being continually relaxed

[37]Fang et al. (2012) estimate that the policy generated a 0.8 year increase in the average number of years spent in middle school. Though we can replicate their results using their chosen empirical specification, when we apply the treatment years identified in that study to our data using our RD design, we find no significant difference between treated individuals in terms of either years spent in middle school or the likelihood that individuals completed at least a middle school degree.

both at the university and high school levels over the period we study (Ministry of Education of the People's Republic of China, 2013).

Supply constraints are most salient at the transition from high school to tertiary education. Among those bunching at the high school credential attainment year, demand exceeding supply of university spots can explain only a fraction of the bunching we observe. In our estimation sample, 56% of high school finishers did not ascend to tertiary education. It is highly unlikely that these individuals finish their schooling in this year because the ultimate year of schooling generates a valuable increase in human capital. In China, the final year of high school is spent almost exclusively on review for the entrance examination to the next level. This year of review, similar to "cram school," is unlikely to contribute to human capital accumulation to the same extent as earlier years in the credential (Larmer, 2014). Another possibility is that the expectation or "lottery" value of sitting the college entrance exam is great enough to make it worthwhile for many to endure the final, review year of high school (Banerjee and Duflo, 2011). National statistics show that, during our study period, 20-50% of individuals who make it to the final year of high school but do not ascend to tertiary education do not attempt the college entrance exam[38]. This means that at least half of the bunching we observe at the end of high school cannot be explained by the lottery value story.

In China, the cultural value of credential attainment is important. This value, however, is almost exclusively placed on receipt of a college degree (Kipnis, 2011; Loyalka et al., 2013). As shown in Figure IV, only a small portion of our sample goes on to this stage and most of the rest know early on if they are likely to ascend to that level. This candidate explanation is therefore of little help in explaining the large amount of bunching at non-tertiary credential attainment years.

The last concern we engage is that measurement error from recall bias leads us to overestimate the extent of bunching. We provide three pieces of evidence that this is unlikely to be the case. First, we note that our CFPS data collects the years in which individuals enter and leave each level of schooling, reducing the risk of reporting bias relative to surveys which ask only how many years are spent in a given level. Second, we see that respondents appear to accurately report when the number of years of primary school needed to attain a credential changes from five to six, even when it is 25 years prior to their being interviewed. This suggests that our data on self-reports of the amount of time spent in each level of schooling are unlikely to suffer from recall bias. Third, we examine drop-out rates in a sub-sample of the educational gazetteers which report statistics on how many children are in each level of school. Overall, we find very low drop-out rates for middle and high school (less than 5% per year), particularly in urban areas, consistent with our data on self-reports of years spent in school.

APPENDIX 4 - DETAILS OF COST-BENEFIT ANALYSIS

This section outlines our data sources and extrapolation assumptions for the cost benefit analysis we perform. The costs of the policy we include consist of one year of forgone wages for all affected individuals. We assume the value of one individual's forgone year of wages is equal

[38] Source: China Education Statistical Yearbooks. This 20-50% estimate masks the fact that few technical high school students take the college entrance exam, while a much larger proportion of regular high school students do so. We are unable to differentiate between those in technical and regular high school in the census data.

to $\frac{\alpha * GDP_t}{N_t}$, where N_t is the size of the active workforce in year t. We calculate if and when the forgone year falls for each individual born on or after 1969 (as the first affected cohort would be leaving primary school in 1981). To calculate the total number losing a given year in each cohort, we multiply the total number of individuals in each cohort (generated using data from the Chinese National Statistical Bureau and World Bank World Development Indicators data) by the proportion of individuals with a given credential in each cohort (calculated from the CLDS data). We assume the lost year is the year a person would have entered the labor force in the absence of the policy, e.g. for someone who was born in 1980 and finishes school with a middle school credential, we calculate the value of the work they would have done in 1995 when they would have left school in the absence of the policy (starting school at age 7, spending five years in primary school and three years in middle school). We assume that those leaving school after primary school do not lose a year of work, as they will leave school at the age of 12 and are not likely to immediately enter gainful employment. For simplicity, we assume that everyone enters school at age 7.

The benefits from the policy consist of the boost in earnings we estimate for affected individuals accruing throughout their lives. This parameter is taken directly from our estimates in Section V. We assume that the benefits affected individuals enjoy are constant throughout their lifetimes, and calculate the total value of the earnings gain for each affected individual up to the year 2050. Table A.6 provides our data sources for each of these figures and assumptions used to extrapolate into the future and where data is not available.

We exclude a few potential sources of costs and benefits, believing them to be several orders of magnitude smaller than the lost year of wages and the lifetime of income gains. On the costs side, we exclude the cost of hiring new teachers and the cost of building new facilities. We do so because archival records suggest that in the short term, the current staff was used to meet most of the increased personnel needs imposed by the policy and classrooms could be split, thus requiring no sizable infrastructure additions. In the long term, the one child policy reduced the number of children in schools, and so existing staff levels and infrastructure house a smaller and smaller number of children. We assume that this would largely offset the need for long term staff and infrastructure adjustments to accommodate the extra cohort of primary students. On the benefits side, we disregard the possible benefit of the creation of new teacher jobs, as our assumption about the staffing costs of the policy implies very few additional jobs would be created. As in Duflo's approach, we exclude the possible beneficial impact of the extra year of schooling on health, fertility and other non-financial outcomes.

Table A.6: Cost-benefit calculation details

Data point	Source	Range of data	Range for extrapolation	Assumptions used	What used for	Notes
Birth rate per 1,000 women	stats.gov.cn	1966-2013	N/A	N/A	Size of cohort	Assume: cohort size constant 2013-2050
Number of women	stats.gov.cn	1970-2013	N/A	N/A	Size of cohort	-
Total labor force	World Bank WDI (World Dev. Indicators)	1981-2015, five-yearly estimates for 2020-2050	Gaps in estimates for 2020-2050	Linear trend between estimates	Total labor force	Close to stats.gov.cn figures, but with predictions to 2050
Proportion of cohort with each education level	China Labor-force Dynamics Survey (CLDS)	Birth cohorts 1969-1987	Birth cohorts 1987-2035	Linear trends from previous 9 years and rules of probability ($0 \leq P \leq 1$)	CB estimates with cohort-specific effects	Begin extrapolation at 1987 cohort to avoid measurement error
Proportion of cohort affected by policy	China's National Bureau of Statistics	1981-2010	2010-2050	Full coverage after 2010	Proportion affected by policy	Equate 5 year primary + 4 year middle with 6 primary + 3 middle
Alpha (labor's share of GDP)	Karabarbounis and Neiman (2014)	1992-2009	1981-1991 2010-2050	1981-91: average of t+1 to t+3; 2010-50: average of t-1 to t-3	Argument in final calculation	-
GDP (constant 2005 US$)	World Bank WDI	1981-2013	2014-2050	Growth rate declines annually from 2013 rate to 4%, annual decrease of 0.15%	Argument in final calculation	-