Demand for Schooling, Returns to Schooling, and the Role of Credentials

Alex Eble and Feng Hu*

July 2015

Abstract

Wages are positively correlated with years of schooling. This correlation is largely driven by two mechanisms: signaling and skill acquisition. We exploit a policy change in China to evaluate their relative importance. The policy, rolled out from 1980 to 2005, extended primary school by one year. Affected individuals must then complete more schooling to obtain their highest credential, the main signal of interest. If the primary mechanism behind schooling returns is signaling, we would expect little change in the distribution of credentials in the population, but a large increase in schooling. If skill acquisition dominates, we should see no change in length of schooling but a change in credentials. Our results are consistent with the signaling story. Further consistent with such a story, we estimate that the labor market return to another year of schooling is very small, though greater for the less-educated. We estimate that this policy, while redistributive, generates a likely net loss of at least tens of billions of dollars, reallocating nearly one trillion person-hours from the labor market to schooling with meager overall returns.

*Eble: Department of Economics, Brown University. 64 Waterman Street, Providence, RI 02912. Email: alexander_eble@brown.edu; Hu: Dongling School of Economics and Management, University of Science and Technology Beijing, 30 Xueyuan Road, Haidian District, Beijing, China 100083. Email: feng3hu@gmail.com.
We would like to thank Anna Aizer, Marianna Battaglia, Nate Baum-Snow, Ken Chay, Andrew Foster, David Glancy, Nate Hilger, Rob Jensen, Melanie Khamis, Sri Nagavarupu, Gareth Olds, Emily Oster, Anja Sautmann, Rajiv Sethi, Jesse Shapiro, Zach Sullivan, John Tyler, and seminar participants at Brown, NEUDC 2013, Wesleyan, APPAM 2014, and SOLE 2015 for helpful suggestions. Eble gratefully acknowledges support from the United States National Science Foundation through a Graduate Research Fellowship and an IGERT Fellowship, and financial and computing resources from the Brown University PSTC. All remaining errors are our own. JEL1: O15, JEL2: I25

1 Introduction

Wages and years of schooling are strongly positively correlated both in developed and developing countries. The cross-sectional estimates of the average labor market premium to an additional year of schooling range from 10-30% (Psacharopoulos and Patrinos, 2004). Economists believe that these returns largely flow through two channels (Weiss, 1995). The first is that schooling endows students with skills that are valuable to employers ("skill acquisition"). The second is that schooling acts as a sieve, putting students through a costly process which provides employers with information about the level of a valuable but unobservable endowment ("signaling"). This information is most often conveyed through an individual's highest educational credential, for example a college diploma (Hungerford and Solon, 1987).

The ideal policy experiment to tease apart the relative contribution of these two channels would hold either highest educational credential or years of schooling constant while varying the other. One option is to vary whether a credential is conferred, holding constant years of schooling, as in Tyler et al. (2000) and Clark and Martorell (2014). We could also study the response to a policy which varies the amount of schooling it takes to get a credential. If skill acquisition is the dominant source of the returns to schooling, this type of policy should keep years of schooling constant and therefore change the share of individuals with a given credential. If signaling is the main source of returns, such a policy will see little change in credentials but a change in years spent in school. In addition, if signaling is the primary source of returns, we should also see that the marginal change in skills induced by the policy should have limited returns. In this paper, we study a policy change in China which hews very closely to this second type of experiment.

In 1980 the Chinese government announced that it would increase by one the number of years needed to complete primary education while leaving unchanged the national curriculum and length of all other levels of schooling. This policy was rolled out gradually over 25 years and has induced over 400 million individuals to spend an additional year in primary school so far. The policy was implemented in each locale at a time when over 75% of individuals proceeded to middle school or beyond and the median student left school after receiving her middle school diploma. When affected by the pol-

icy, this median student could either keep total years of schooling constant by forgoing the credential she would have attained in the absence of the policy[1], or keep credential attainment constant by spending an additional year in school. Figure 1 gives a graphical representation of the Chinese education system pre- and post-policy as well as the two models' predictions of how the median student would respond.

We identify the causal effect of the policy on schooling and labor market outcomes using a simple regression discontinuity design. We compare treated and untreated individuals within each locale, restricting our attention to those leaving primary school within a few years of when the policy takes effect. This approach is similar to recent work studying the impact of changes in compulsory education in the UK (Oreopoulos, 2006; Clark and Royer, 2013). Unlike these studies, which examine the impacts of a policy implemented at once across an entire nation, we take advantage of the fact that our policy was rolled out at different times in different places across China to net out cohort, place, and cohort-by-region fixed effects. This protects against the threat of upward bias from geographically autocorrelated secular trends, shown to affect similar exercises in the US (Stephens and Yang, 2014). To determine the treatment status of individuals in our data, we conduct a large archival data collection exercise, examining records from each of China's 345 prefectures documenting if, when, and how the policy was implemented.

We begin our analysis by estimating the impact of this change on credential attainment and years of schooling. As noted above, if schooling returns are primarily due to signaling, optimal response to such a policy should increase schooling (by approximately one year) and leave credentials unchanged. If skill acquisition is the primary source of returns, we should observe a constant number of years of school and a change in credentials. Our results are consistent with the signaling model's predictions. We find evidence that years of schooling increased by nearly one for affected individuals, and there was no evidence of a decrease in credential attainment or years of post-primary schooling to offset the extra primary year. This means the policy has reallocated nearly a trillion person-hours from the labor market to the pursuit of schooling. We also find no evidence that the policy changed the characteristics of who gets what credential, which allows us

[1] Though middle school was made compulsory in 1986, in Section 6.2 we show evidence that the rollout of the policy has little effect on whether or not individuals complete middle school or earn at least a middle school credential.

to estimate the return to an extra year of primary school, holding highest educational credential constant.

Our estimates of the returns to schooling also suggest the importance of the signaling channel. We exploit the large sample size of the 2005 Chinese by-census to generate a precise but small estimate of the labor market returns to this additional year of primary school, holding highest educational credential constant. 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 channel and not the result of the policy influencing wages through selection into different types of employment. These returns are quite small relative to the 8-14% per-year premium to earning a credential seen in these same data. The extra year of primary school is review, but this is similar to the last year of middle and high school, both of which also devote a large portion of instructional time to review for entrance exams. These results suggest that more than half of the per-year returns to earning a credential, naively estimated in the cross section, are likely flowing through the signaling channel.

We estimate higher returns from this extra year accruing to those with less schooling, the intended beneficiaries of the policy. This result is akin to recent work showing that increased instructional time is an important mechanism behind the success of a group of high-performing New York charter schools in improving outcomes for the underprivileged (Dobbie and Fryer, 2013). We measure no difference between the treated and untreated in cognitive ability scores, either overall or at several quantiles of interest. This adds to recent evidence that there can be measurable adult labor market returns to childhood interventions long after any measured cognitive ability gains wane (Heckman, 2006; Chetty et al., 2011).

Both the response to the policy and the estimate of the returns to the additional year of schooling suggest that a large portion of the returns to schooling in China flow through the signaling mechanism. We develop a simple model of the household's schooling decision to formalize the reasoning behind our interpretation and discuss the plausibility of the assumptions needed to apply its predictions to our data.

Our paper contributes new evidence to two lines of empirical work on signaling and skill acquisition. One set of studies highlights a set of cases where schooling decisions are driven by the signaling value of attaining a credential (Lang and Kropp, 1986; Bedard, 2001). To it, we contribute evidence from a more general case where the entire population is shocked with a change in the length of time required to obtain a credential. The second set of studies shows significant positive returns to endowing an individual with a signal while holding years of schooling constant. This work uses naturally occurring, plausibly exogenous variation to causally estimate the value of a signal, exploiting either arbitrary test passing thresholds or comparing otherwise similar education reforms that vary in whether they also confer a signal on the treated population (Tyler et al., 2000; Grenet, 2013; Clark and Martorell, 2014). Our study contributes evidence from the converse of what these papers study - we vary years of schooling, holding highest educational credential constant, and estimate the returns to this marginal change in human capital.

Though our estimate of the return to an additional year of schooling is 3-7 times smaller than most credibly identified estimates from the developed world (Angrist and Krueger, 1991; Oreopoulos, 2006), it is similar in magnitude to recent work using twins in China (Li et al., 2012b). Our estimate is also larger than the zero estimated effect on earnings of extending the length of the basic track of German high school by one year (Pischke and von Wachter, 2008).

A primary impact of this policy is the reallocation of 0.85 trillion person-hours from work to school. We perform a cost-benefit analysis exercise leaning heavily on that of Duflo (2001), comparing the value of the increase in income gained from the extra year to the lost year of productive activity in the labor market. As in Duflo's study, our time frame ranges from the first graduating cohort affected by the policy in 1981 to 2050. This exercise generates four sets of estimates. In all but the most favorable scenario, the costs of the policy exceed its benefits by at least tens of billions of 2005 US dollars.

Finally, we draw policy implications from our research. We estimate that the policy has moved nearly one trillion person-hours[2] from the labor market to schooling so far, at a likely fiscal loss. This result highlights the importance of a seemingly arbitrary policy

[2]National statistics on rollout of the policy indicate that more than 400 million individuals have been affected to date. We assume individuals spend 2,000 hours working in a year.

choice: how long should each level of schooling last? We use DHS data to show that bunching at credential attainment years is a relatively common phenomenon in developing countries, and argue that this policy decision is under-researched given the massive resource stakes involved.

The rest of the paper proceeds as follows. In Section 2, we discuss the history of education in modern China and describe the policy we study. In Section 3, we describe the data we use and our identification strategy. Section 4 contains empirical results related to educational attainment and Section 5 provides our empirical results relating to the labor market. In Section 6, we interpret our results using a simple conceptual framework and several ancillary analyses and perform the cost-benefit analysis. Section 7 concludes.

2 A brief history of primary and secondary education in modern China

China's education system has grown substantially in size and scope since independence. In 1945 when the country was founded, education levels were quite low – only 20% of the population was literate, and less than 40% of school-aged children were in school (Hannum, 1999). The country poured resources into primary education after independence, shortening the length of primary school from six years, which it had been traditionally, to five, and vastly expanded the number of schools across the country at all levels (Liu, 1993). In the 1950's, 60's, and 70's, a series of policy experiments and disasters racked China, leaving in its wake an educational system which varied greatly in length and structure across provinces. Nonetheless, literacy rose to more than 50% by 1976 and the average educational attainment of the populace 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[3] in all provinces.

[3] 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.

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 should gradually be extended to twelve years. It encouraged gradual adoption of six year primary school throughout the country, putting more pressure on urban schools (Liu, 1993). In practice, the policy was implemented at different times in different places. The policy decrees that the length change be implemented "according to local conditions," and in hundreds of counties primary education remained at five years of length until the early 2000s. Appendix Figure A.1 plots national data on the proportion of students in six year (or equivalent) primary school systems, showing gradual adoption of six year primary education between 1980 and 2010. About 60 percent of locales switched to a six year system between 1981 and 1993, relatively few made the change in the mid-1990's, and the rest shifted in the late 1990's and 2000's, reaching near-universal adoption in 2007.

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, it was implemented as part of Deng Xiaoping's move in the late 70's and early 80's to prepare China's labor force to "face modernity" (Vogel, 2011). In qualitative interviews that we conducted with parents, teachers, and students who experienced the change, many added that the policy was most likely to benefit those with the lowest ability.

The transition from five to six years of primary education could be done in a number of manners. Appendix Table A.1 gives six examples of how the policy was enacted, taken from gazetteers in different implementing cities and counties across the country. In some cases, this 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 3, most counties had ultimate say on the year in which the switch was made. Our data bears this out, and in later sections we address the issues surrounding this discretionary 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 education from five to six: by 1984, 90% of students in high school were in three year programs; in contrast, it wasn't until 2003 that more than 90% of primary school students were in six year programs (National Institute, 1984).

3 Data and identification strategy

This section describes the data sources and empirical methods of the paper. We match archival data on policy implementation to data from the census and large, nationally representative cross-sectional surveys to estimate the causal impact of the policy on educational and labor market outcomes using a regression discontinuity design. 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.

3.1 Data sources

We use three main sources of data, listed in Table 1. The first is the 2005 Chinese by-census, which collects basic data on family structure, highest educational credential attained, health, and income, and contains 2.6 million observations[4]. The second is the China Family Panel Studies (CFPS), a new, nationally representative dataset containing

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

information from over 30,000 individuals in rural and urban China across 25 provinces, representative of 94.5% of China's population[5]. The third dataset is the China Labor-force Dynamics Survey, which is similar to CFPS in its scope and ambition, but focuses on participants' work and other economic activities. The latter two data sets collect detailed information on how many years individuals spent in each level of schooling (e.g. five in primary, three in middle school, and one in high school), which allow us to examine whether the policy we study affected the number of years spent in each level of post-primary schooling as well as highest credential attained.

The policy we study was implemented at different times in different places both across and within China's provinces, as shown in Figure 2. We hired a team of research assistants to comb through county educational gazetteers[6] stored in the Chinese National and Peking University Archives, to determine if, when, and how the policy was implemented in each locale, as in recent work by Almond et al. (2013). We are able to determine the year the policy was implemented in 280 of the 345 prefectures[7] 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 prefecture-level treatment year, or 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.

3.1.1 Determining treatment status from the data when geographical data are anonymized

The CFPS dataset has a few traits which make it particularly desirable - as in the Labor-force Dynamics Survey, it collects detailed data on how many years individuals spend in each level of schooling, but its sample size is twice as large and it identifies which indi-

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

[6]Appendix Figure A.4 shows a page from one of these gazetteers.

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

viduals 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.

CFPS county names are anonymized, so we cannot use the archival records to assign treatment status. Instead, we apply an algorithm to the data to generate, for each county, the most likely cohort in which the number of years spent in primary school jumps from five to six. The mean shift model we apply (Fukunaga and Hostetler, 1975) is similar to that used in Munshi and Rosenzweig (2013). It is a simple statistical tool, often used in fields such as machine learning and digital recognition, to identify where structural breaks are most likely in a dataset. Its implementation in our context is straightforward - separately by county, we regress individual-level years of primary education on a constant and an indicator function for having graduated in or after a given year:

$$s_i = \beta_0 + \beta_1 * \mathbf{1}\{t_i \geq t^*\} + \epsilon_i \tag{1}$$

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 of primary education for individual i, t_i is the year in which she graduated from primary school, and ϵ_t 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[8]. This exercise generates a treatment year for each county in our estimation sample[9]. In Appendix A3, we use national statistics and the application of both archival and mean shift methods to the China Labor-Force Dynamics data to corroborate the reliability of the mean shift method's identified treatment years.

[8]An example of this process is shown in Appendix Figure A.5, which shows the histogram of mean years of primary education in a county and plots the ssr estimates generated by Equation 1 for each treatment year. The ssr sequence reaches its global minimum at 1997, where we also observe a clear shift upwards in mean years of primary education from approximately five to six.

[9]The CFPS collects data from 144 counties across China, excluding Shanghai, which implemented the policy by extending the length of middle school instead of primary school. We include 112 of these 144 counties in our final estimation sample. Appendix A3 lists the inclusion criteria used to determine this sample. Our empirical results are robust to using data from all 144 counties.

3.2 Empirical strategy

Our identification strategy is a simple regression discontinuity design. We compare outcomes of individuals finishing primary school just before the policy is implemented in a given county to those in the same county 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[10].

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 urban residence permit holders. Figure 3 plots these data, condensed to distance to treatment year means, and 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 Lee and Lemieux (2010) and Imbens and Lemieux (2008), our main estimating equation is an ordinary least squares regression of y_{pi}, the outcome of interest for individual i in place p (either county or prefecture), on a short set of key regressors:

$$y_{pi} = \beta_0 + \beta_1 * T_{pi} + \beta_2 * f(t_{pi} - t_p^*) + \beta_3 V_i + \lambda_t + \mu_p + \eta_{tr} + \epsilon_{pi} \qquad (2)$$

Here T_{pi} is an indicator variable equal to 1 if the individual, i, belongs to a cohort finishing primary school in or after the identified treatment year in her place, p. The term $f(t_{pi} - t_p^*)$ is a linear function of the time elapsed between the place-specific year of treatment, t_p^*, and the year in which the individual graduated from primary school, t_{pi}. $f(\cdot)$ is linear

[10]Estimation equations such as ours without cohort-by-region fixed effects are shown in Stephens and Yang (2014) to generate upwardly biased estimates.

and is estimated separately[11] for treated and untreated groups (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. Place (μ_p), cohort (λ_t), and cohort-by-region (η_{tr}) 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 our two main empirical results. For the rest, we estimate on 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 et al., 2004). We restrict our time frame to cohorts leaving primary school between 1976[12] and several years before the sample is drawn to give most individuals enough time to finish their educational career before being observed (1995 in the census data, and in 2003 in the CFPS data).

3.3 Potential confounders

The implementation of this policy across space and time was decided upon by provincial and county-level bureaucrats. We have shown evidence that our main identification assumption is upheld and, as we are identifying off of differences within each county, this local discretion is potentially of less concern for us. It may, however, raise specific concerns of omitted variable bias which must be addressed for causal interpretation of results.

If there were another policy or external phenomenon correlated with both when the policy was implemented in a given place and the later educational decisions or labor market outcomes of affected individuals, our results would suffer from omitted variable bias. To address this concern, we take a few measures. First, as recommended by Stephens and Yang (2014), we control for cohort-by-region fixed effects. If such an exter-

[11] We include this to capture two possibilities: one, that the effect may differ over time elapsed since the treatment year, as counties get better at implementing the policy, and two, that there may be an existing pre- or post-trend that we wish to control for so that β_1 captures only the difference between pre- and post-policy means. In Appendix A.1 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.

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

nal phenomenon were geographically autocorrelated, 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 2 provides a heat map of prefecture implementation years, with lighter shades indicating earlier implementation. This map shows a wide distribution of timing of implementation with no obvious geographical pattern beyond later implementation in some provinces 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. Ultimately, it is impossible to conclusively rule out existence of an unknown omitted factor driving both the local bureaucrat's choice of when to implement the policy and individual educational decisions and later labor market outcomes. We have shown, however, that we are unable to find evidence of one among the set of likely candidates.

A final concern is the potential for migration to bias our results. Despite the best efforts of the data collectors, we are less likely to observe migrants than non-migrants. If the treatment effect is different for migrants and non-migrants or the treatment affects who is likely to migrate, our estimate of the treatment effect will differ from the 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 the effect of the treatment on size and characteristics of individuals in our rural sample, and find no significant relationship between treatment status and cohort size or cohort gender composition. Beyond these tests, there is not much 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 CFPS reached (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).

4 Empirical results - education

In this section, we estimate the impact of extending the length of primary school by one year on individuals' later educational outcomes. First, we show that the policy was indeed

effective at increasing the number of years individuals spent in primary school. We then show our treatment effect estimates for years spent in post-primary education, 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 the characteristics of individuals with each credential.

This policy was implemented in each county at a time when almost 80% of students went on to get at least some post-primary education. If human capital accumulation were the main goal of attending school, these individuals would be likely to hold total years of schooling constant by offsetting the additional year of primary school with one less year of post-primary school[13].

We first examine whether the policy had its desired effect of increasing primary school for affected individuals. Figure 4 plots distance-to-treatment year means of the proportion of individuals getting six or more years of primary education in our CFPS sample overlaid on estimates of their confidence intervals. Prior to implementation of the policy, the proportion of students getting at least six years of primary school is between 20 and 30% of the population. At the policy implementation year it jumps to over 80%, increasing to nearly 100% in the decade following implementation. Results from the regression analog to this exercise are presented in the first row of Table 2. We estimate 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.

Figure 5 shows, separately for untreated and treated observations, the PDF of the number of years of post-primary education. This figure also uses CFPS data, and restricts the sample to students who graduate from primary school within five years of the policy implementation year. This figure nicely summarizes our main empirical results related to the effect of the policy on post-primary education. We see extensive bunching at credential attainment years[14], and little visible difference between the treated and untreated groups in either the location or the magnitude of this bunching, suggesting that there was no net offsetting behavior in response to the lengthening of primary school.

[13] In the absence of the extra year of primary school conferring a major gain in skills which induces individuals to proceed further in schooling.

[14] 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.

The regression results for our education outcomes are given in the rest of Table 2. The second row shows our estimate of the effect of the policy on total years of education to be 0.660, significant at the 1% level. This result is consistent with all individuals induced by the policy to spend an additional year in primary school displaying no offsetting downward adjustment in post-primary schooling. 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 6, we show that our point estimate on post-primary education 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[15].

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 2, we estimate the effect of the policy on whether or not an individual gets at least a middle school credential and whether or not she gets at least a high school credential using both census and CFPS data. The effect of the policy on middle school completion estimated on the census data is negative but small (0.0049, from an untreated group average of of 0.725) and insignificant (SE 0.0030), and we can reject anything larger 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 not significant at conventional levels. 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. We interpret the totality of these results as evidence there was no systematic downward adjustment of post-primary schooling to offset the additional year of primary school, consistent with the signaling model's predictions.

[15] In Appendix Figure A.7, we plot the trends over distance-to-treatment year bin means in raw data (left column) and cohort-demeaned residuals (right column) for total years of education (top row) and years of post-primary education (bottom row). We see the same patterns as in the regression coefficients: an upward jump of about one year of total education at the treatment threshold, and no downward jump in post-primary education at the treatment threshold, both for the raw and demeaned data.

One possible explanation for this overall pattern of no net change in post-primary schooling is a change in composition of who gets 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 at each level of schooling.

The subgroups we expect to be most affected by the policy are women, as previous work has shown that Chinese households in the 80's and 90's often chose to allocate fewer resources to women (Li, 2003); rural residents, because of the lower income levels in rural areas; and women from rural areas, who are doubly disadvantaged. Table 3 shows the subgroup-specific treatment effect estimates for the same educational outcomes examined in Table 2. The estimated treatment effects are largely negative, as predicted, and consistently so for rural women, the most disadvantaged group. 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). Our set of predetermined characteristics serve as the outcome variables. Wald tests of the equality of the treatment-by-credential level coefficients tell us whether the proportion of individuals with a given set of predetermined characteristics within a credential changes relative to that proportion in the other credentials across the treatment threshold. We use CFPS and CLDS data for these tests, and fail to reject equality of the treatment-by-credential coefficients on an individual's number of siblings, mother's and father's education, and gender (p-values 0.35, 0.93, 0.83, and

0.66, respectively).

5 The labor market

In this section we estimate the impact of the policy on various labor market outcomes, including employment status, type of employment for the employed (i.e. entrepreneurship, 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 capitalist labor markets as early as the late 1990's (Cai et al., 2008). We use the 2005 by-census data for all of the analyses in this section for its sample size, and in 2005 we expect most workers to earn wages that are at least strongly correlated with their marginal product (Zhang et al., 2005). 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 loss to migration is more of a concern. From our results on whether or not the composition of individuals in each credential level changed, we conclude that the main effect of the policy was to force nearly everyone to spend an extra year in primary school while holding their highest credential constant. Under that assumption, we can use our identification strategy to isolate the labor market returns to this extra year of school from the confounding effect of receiving an additional credential that often biases such work (Weiss, 1995).

Our main dependent variable of interest is the natural logarithm of monthly income. 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 whether she is an entrepreneur. In the regression results presented in this section, we add educational 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.

Regression results are given in Table 4. 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. We conclude from this that it is unlikely that the policy has its impact through a selection of treated individuals into different jobs.

Next, we estimate the effect of the policy on the natural log of monthly income. We add employer-type fixed effects to control for changes in China over time between professions. Our first specification uses cohort and place (prefecture) fixed effects. Here we find a gain of 1.94% in monthly income, statistically significant at the 99% level. Recent work by 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 region-by-cohort fixed effects to mitigate this bias. We next execute our specification with cohort fixed effects specific to each of China's regions (East, Northeast, Central, and Western). 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 6 presents our treatment effect estimates for log monthly income and their confidence intervals for the same nine different choices of sample bandwidth. This figure demonstrates the stability of both the magnitude of the coefficient and its ability to reject a zero effect across a wide range of bandwidth choice.

There is still concern that, given the secular increase in education levels and incomes in China over this period, our research design comparing those graduating from primary school earlier to those graduating later may be "stacking the deck" in favor of finding a positive and significant estimate. To test for this, we conduct a Monte Carlo exercise in which we draw 1,000 placebo treatment years for each county (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 7 gives the probability density function for these estimates. The placebo treatment effect estimates are normally distributed, with a mean of -0.001 and a standard error of 0.0084, putting the true estimate of a 2.03% increase

in wages well beyond two standard deviations from the mean. We conclude that the significance and sign of our estimates are not merely a mechanical result of our research design.

We next explore heterogeneity in treatment by subgroups, shown in the lower rows of Table 4. The measured effect is monotonically decreasing in level of education, consistent with a simple model of diminishing returns to training at a given level (i.e., diminishing returns to reviewing the primary school curriculum) and in accord with the stated intention of the policy to benefit those less well off. As mentioned in the introduction, this result is 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).

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 marginal product of labor in the private sector than in the government (Li et al., 2012a). We observe 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 et al., 2013). Independently run Wald Tests reject equality of coefficient 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 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 place to generate precise estimates using the RD design as specified.

Our average treatment effects are small relative to the benchmark estimates of the returns to an additional year in developed countries (Oreopoulos, 2006). A 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 affects 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[16].

We next test for a difference between the treated and untreated in cognitive ability, as measured by a test administered to adult respondents in the CFPS survey. Figure 8 plots the kernel density functions for treated and untreated individuals using the five year bandwidth sample. The two distributions track each other quite closely. Statistically, we cannot reject their equality using a Kolmogorov-Smirnov test, and quantile regressions run simultaneously at the 20th, 35th, 50th, 65th and 80th quantiles all fail to reject a zero effect of the treatment on cognitive ability. These results notwithstanding, we believe that the earning gains we estimate are likely to be the result of skills acquired in the extra year of primary school. If this is true, our results add to the evidence that childhood interventions which initially generate increases in cognitive ability often bring labor market returns well after measurable ability gains fade (Heckman, 2006; Chetty et al., 2011).

Our estimates of the effect of the policy on monthly income are small compared to naive estimates of what an additional year of education brings based on cross-sectional observation at the time our census data are collected. Table 5 plots the average monthly income for an individual holding each of five credentials - primary, middle, high school, technical college, and university. In column (1), it gives the per-year income bonus of earning that credential, calculated from the credential bonus (shown in column (2)) divided by the number of years it takes to get that credential. Our naive estimates of the per-year premium to earning a credential are between 4 and 7 times as large as our average treatment effect estimate of the return to the extra primary year. Although the sixth year of primary school is a review of the previous five years, the final years of middle school and high school are also largely review in preparation for an entrance exam. This suggests that estimates of the returns to this sixth year are likely to be similar to those to other years of schooling. The discrepancy between the measured impact of the sixth year and the per-year premium of attaining a credential suggests that skill acquisition is

[16] 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 everyone benefitting from 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).

unable to explain the entire per-year credential premium we observe in the cross-section.

We can use these results to make a back-of-the-envelope calculation about the relative contribution of the signaling and human capital channels. If we assume that the final year of middle school and high school, which is also a review, give a similar boost to earnings, and other years give twice that boost (i.e. 4.06%, slightly larger than the per-year estimate Li et al. (2012b) generate using twins for identification), then the signaling channel would account for 57.2% of the returns to a middle school degree and 60.0% of the returns to a high school degree.

6 Interpretation and discussion

This section of the paper has three parts. First, we model a household's decision about their child's schooling to identify conditions under which we can detect the signaling motive vis a vis skill acquisition in our empirical results. We then enumerate a set of other mechanisms which could generate similar results, describing their empirical plausibility in our setting. We finish the section with a cost-benefit calculation of the policy's impacts.

6.1 Modeling the schooling decision

To formalize how we interpret our empirical results through the lens of the signaling/skill acquisition literature, we use a simple model of a household's schooling decision based on Becker (1975). In this model, a household decides after how many years, $s \in [0, S]$, their child leaves (post-primary) schooling[17] and enters the labor market[18]. This decision's arguments are the expected future benefits of schooling, $b(s)$, and the costs of schooling, $c(s)$. We assume that the benefits are increasing and concave, i.e. $b'(s) > 0$ and $b''(s) < 0$, and that costs are increasing and convex, $c'(s) > 0$ and $c''(s) > 0$.

[17] Recall that in our estimation sample, nearly 80% of observations go on to get at least some post-primary schooling and primary school is compulsory for all.

[18] We abstract from discussions about investment in children depending on the bargaining outcome between individual parents and/or parents and the child. For a treatment of these issues, see Bobonis (2009). We also abstract from issues of multiple children and inequality aversion here. As educational levels are rising over this period, multiple children and inequality aversion on the part of the parents will only make it more likely for us to observe a downward shift in post-primary attainment among treated individuals, as the shock of "extra" education to one child raises expenditure and also pushes the parent to reallocate from the affected child to unaffected children. We see no such downward shift. Finally, we assume once a child leaves school, she is unable to return.

The first order condition of this calculus determines the optimal level of schooling, $s*$, satisfying $b'(s*) - c'(s*) \leq b'(s) - c'(s)$ for all $s \neq s*$ and $b'(s*) - c'(s*),\ b'(s) - c'(s) \geq 0$.

We introduce the competing motives of signaling and skill acquisition by making different assumptions about the shape and continuity of $b(s)$, and generate heterogeneity in schooling choice across households by introducing the role of household wealth in the cost of schooling, e.g. $c(s) \to c(s, w)$. Specifically, a family's wealth determines the marginal utility that would be earned from the child leaving school and entering the labor market. The utility cost of a child going to school is negative and decreasing in wealth, $\frac{\partial c(s,w)}{\partial w} < 0$ and $\frac{\partial^2 c(s,w)}{\partial w^2} < 0$, reflecting the notion that the opportunity cost of a child's time is decreasing in the level of consumption her household enjoys[19]. We assume that the distribution of wealth in the population is continuous over a range that produces a non-degenerate distribution of schooling in the population in both the signaling and skill acquisition scenarios.

The first link from this framework to our empirical results has to do with the distribution of schooling in the population, as in Lang and Kropp (1986) and Bedard (2001). In the world where skill acquisition is the only benefit that accrues with schooling, we assume that the acquisition of skills is a continuous, gradual process, $b(s)$, with $b(s)$ smooth and nonzero within a range $(\underline{b}, \bar{b})$ such that the distribution of schooling is not degenerate. Under these assumptions, the distribution of schooling in the population would then closely resemble the distribution of wealth.

Signaling is introduced by adding the assumption that at certain levels of s, $s_i \in \{S_i\}$ a signal is conferred and b'(s) discontinuously increases. These spikes are meant to capture the notion of "credentialism", that obtaining a credential confers benefits in the labor market outside of the process of skill acquisition (Spence, 1973; Weiss, 1995). Adding this assumption changes our prediction of the distribution of schooling in the population for the same distribution of wealth. Formally:

Proposition 1: Let $f(s_l)$ be the probability mass function at schooling level s_l. Under credentialism, $f(s_i) > f(s_j) < f(s_k)$ for two consecutive credential-attainment years[20] $i, k \in \{S_i\}$, and any non-credential attainment year $j \notin \{S_i\}$ satisfying $i < j < k$.

[19] In the Appendix, we develop predictions for a model which adds an additional layer of heterogeneity in ability endowments. As we do not have ability measures before children finish primary school, we are unable to take its predictions to our data.

[20] For example, the year one finishes middle school and the year one finishes high school.

This result follows directly from our assumptions about the marginal benefit of credential and non-credential attainment years. Under our assumptions about the smoothness and regularity of the cost function, the extra benefit conferred by the credential induces bunching at credential attainment years. In the absence of credentialism, the assumptions we have made so far imply no such bunching. Figure A.9 shows the distribution of schooling, assuming a uniform distribution of wealth, for the two regimes.

This framework also generates substantially different predictions of the response to the policy we study. If we assume that the main effect of the policy for the majority of the population is to raise the cost of all additional post-primary schooling by the additional forgone year of wages it requires, we generate Proposition 2:

$$\sum_i \mathbf{1}\{s_{i_{post}} \neq s_{i_{pre}}\} \mid_{credentialism} < \sum_i \mathbf{1}\{s_{i_{post}} \neq s_{i_{pre}}\} \mid_{no\ credentialism}$$

In prose, the skill acquisition model predicts a decrease in the amount of post-primary schooling for a larger proportion of the population than the signaling model, which predicts no change in the equilibrium number of years of post-primary schooling for most, save the few for whom the increase in marginal cost is enough to edge the marginal cost schedule above the marginal benefit peak of the pre-policy equilibrium attainment level. The proof for this is straightforward. The main credentialism assumption is that there exist credential attainment years $s_i \in \{S_i\}$ such that $b'(s_i) > b'(s_j)$ for all $s_j \notin \{S_i\}$. This implies that for a given cost change $c \rightarrow c'$, $P(\frac{\partial s*}{\partial c} = 0) \mid_{credentialism} > P(\frac{\partial s*}{\partial c} = 0) \mid_{no\ credentialism}$ if $s* \in \{S_i\}$ and we have our result. If we relax our previous assumption on the nature of the distribution of wealth, the strong inequality becomes weak. Appendix Figure A.10 displays a case demonstrating with the intuition behind this proposition.

This simple model formalizes how we interpret our empirical results. We see both extensive bunching at credential attainment years and little difference between the pre- and post-policy distribution of schooling, both predictions of the credentialist model. Of course this it not conclusive proof of signaling; the model presented in the previous subsection abstracts from many important aspects of reality which could generate bunching at credential years in the absence of credentialism as we have defined it. In the next

subsection, we list the most likely candidates, describe their theoretical implications, and discuss their empirical plausibility in explaining our results.

6.2 Alternative explanations for bunching at credential attainment years

The main competing hypothesis that could explain our education results is that students were unable to adjust their final level of education either upward or downward, and so the patterns we observe would be best explained by the institutional details of the Chinese educational system and not the signaling value of attaining a credential. With respect to downward adjustment, the first and foremost concern is that China's 1986 compulsory education law, which stipulated that all Chinese students had to complete primary and middle school, may have prevented many students in our study from adjusting downward the number of years of middle school they attained. Recent empirical work has shown the law was gradually implemented and remains porously enforced, with many locales not implementing it until the mid-2000's (Fang et al., 2012). That study shows the policy induced individuals to spend 0.8 extra years in middle school. Though we can replicate their results using their 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.

Furthermore, as recently as 2010, the proportion of students not completing middle school was estimated at over 20% in some rural areas (Mo et al., 2013). In Panel A of Appendix Figure A.8 we plot cohort means of total years of education and superimpose a plot of the cohort-specific proportion of students who get two years of post primary education or less (recall that a middle school credential usually takes three). Though there is a downward trend after 1986, the proportion of students not completing middle school is stable at around 20% throughout most of the period we study, suggesting that a large proportion of students in this period would have been able to adjust on the lower margin should they have desired to do so.

A related possibility is that large increases in costs between schooling levels are generating the bunching at the final year of each credential. There are in fact substan-

tial increases in costs between schooling levels and there is evidence that these form binding constraints for China's rural poor (Liu et al., 2009). For the majority of the population, however, this appears not to be the case. We know that 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 ascent to higher levels of schooling is oversubscribed (Kipnis, 2011; Butrymowicz, 2012; Larmer, 2014) While there is no conclusive study on the matter that we are aware of, we interpret the available evidence as indicative that the discrete jumps in costs between levels of schooling in China are unlikely to explain the majority of the bunching we observe, particularly not in urban areas where incomes are higher and costs are less likely to be a constraining factor

This pattern of oversubscription speaks to the other side of the institutional rigidity argument - supply constraints. Demand for schooling exceeding supply could also generate the bunching and non-response of post-primary schooling to the extension of primary school we observe. Though we know that the number of college entrance exam takers has greatly exceeded the number of college entrance spots for the last half of the 20th century, national-level data shown in Panel B of Appendix Figure A.8 provides suggestive evidence that over the period we study supply constraints are being continually relaxed both at the university and high school levels (Ministry of Education of the People's Republic of China, 2013).

Even in the presence of these supply constraints, we are left with the question of why so many individuals finish school at the credential attainment year. In our estimation sample, 54% of middle school finishers did not ascend to high school, and 56% of high school finishers did not ascend to tertiary education. Ideally, we would like to know to what extent these individuals, and those who go on to complete university, remain in schooling until the end of a credential because they value what they learn in school as opposed to doing so because they want the signal (or some other benefit) conferred by obtaining a credential. Here we provide suggestive evidence about the likely motives, but a conclusive answer to the question is beyond the scope of this research.

In China, the final year of middle school and high school is spent primarily on review for the entrance examination to the next level. The content of this year is unlikely to be

new or, for that matter, useful for much beyond passing the entrance exam to proceed to the next level. This structure therefore makes it unlikely that skill acquisition alone would explain the bunching at credential attainment years and zero effect of the policy on post-primary schooling that we observe.

It is possible that students who remain in school until the end of a credential do so because they want to advance to the next level of schooling (either for the skills gained there or the signal), and the expectation value of sitting the entrance exam is great enough to make it worthwhile to endure the review year. Again, institutional features suggest this explanation is unlikely to account for the majority of our results. In most schools in China, a student's performance is frequently evaluated, in some cases to an excessive degree (Kipnis, 2011). While there is some evidence that in China and elsewhere in the developing world, households' views on the returns to schooling (Loyalka et al., 2013) and their own children's ability (Dizon-Ross, 2014) can be over-optimistic, the same research on ability also shows that frequent contact with accurate assessments of a child's ability are likely to correct these misperceptions, at least in the context of the household education decision we study[21] (Dizon-Ross, 2014). We argue, then, that a large proportion of these students who finish a credential but do not ascend to high school are not marginal cases for passing the entrance exam[22] and so the "lottery value" concept appears unable to account for more than a small portion of the results we are trying to explain, given what we know about the Chinese system and individuals' response to information about their own ability.

Finally, it is possible that some of the bunching we observe could be merely the result of measurement error from recall bias. If individuals respond that the length of time they spent in a given level of schooling was however many years it takes to finish that level of schooling, regardless of whether or not they spent that many years in that level, we would overestimate the extent of bunching. We provide two pieces of evidence that this is unlikely to be the case. One, we see that respondents appear to accurately report when the years spent in primary school went from five to six, even when it is 25 years

[21] Mobius et al. (2011) study a similar setting the lab, where decision-makers are treated with information that contradicts the decision-maker's perception about his own ability, and find the opposite conclusion, that new information can reinforce a bad prior.

[22] This is impossible to know without access the distribution of publicly announced grades of students in the first and second years of middle and high schools, but our best guess is that more than 50% know that they will not ascend.

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. Second, we look for 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.

6.3 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, 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[23], w_{it}, that affected students i forgo during the year t they spend in school instead of in the labor market:

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

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[24]. We then calculate the total value of the value of 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

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

[24] For example, an individual born in 1975 would start school in 1981. If their highest credential attained is a middle school degree (and so spent nine years in school) and they were affected by the policy, their "lost year" would be in 1990.

is no productive activity displaced by the policy other than the reallocation of students' time in the workforce.

The benefit of this policy is estimated to be the sum of all wage gains for all affected cohorts over the time frame we have chosen:

$$Benefit = \sum_{t}\sum_{i} \alpha GDP(t)S(c,t)P(c,t)\beta$$

Here α is the share of labor in GDP[25], $S(c,t)$ is the fraction of total wages earned by cohort c 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 wages. 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[26].

In Table 6, we present four estimates, varying two important 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 schooling in each cohort. 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, differential switching between agriculture and the non-farm labor market, and the great changes affecting rural China over this period. The policy, however, was implemented in both urban and rural China. Furthermore, until 2011, more than half of China's population was rural. We present two sets of estimates - in the first row, we use the urban estimates on returns for urban and rural residents. In the second row, we use the treatment effect estimates for rural areas (first the average treatment effect, then the credential-specific effects estimated off of only rural residents).

We give both the net and gross benefit of the policy over this period under each of the four scenarios (type of treatment effect by treatment of rural areas). Under three of the four scenarios, the cost of the treatment exceeds the benefit in wages by tens of billions

[25]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.

[26]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 in which benefits were enjoyed and thus reduce our estimate of the benefits of the policy.

of dollars. In the scenario in which we impute the urban, credential-specific estimates to rural residents, we see a positive estimate of about four billion dollars. Furthermore, 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. Such costs are likely to 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.

7 Discussion and policy implications

In this paper, we exploited a massive policy change in China's educational system to study how household decisions on children's post-primary schooling respond to extending the length of primary school by one year. We find that the vast majority of individuals chose to spend another year of school and transport fees and forgo a year of wages to get the credential they would have attained in the absence of the policy. We used the results from this exercise to generate a new parameter estimate of the returns to a year of education, holding highest educational credential constant. The extra year of school generates a small but precisely estimated two percent increase in monthly income which is higher for those with less education, those who the policy set out to assist.

We then framed these results in the context of the literature on the contributions of signaling and skill acquisition to the returns to schooling. We interpret our results as suggestive that the signaling value of attaining a credential is the dominant driver behind schooling decisions in China in the modern era. 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 education, finding that in most scenarios the policy will be a net loss of tens of billions of 2005 US dollars.

We conclude 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 Demographic and Health Surveys data from those 74 countries in which the number of years a respondent

spends in school is collected. In 48 of these countries[27], we found evidence of bunching at multiple years similar to that shown in Figure 5. 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 the decisions on the length of credentials have massive resource implications.

Though the extra year appeared to be a net loss in China's case, it was redistributive: those with lower levels of education gained the most from it. In many countries there appears to be similar bunching at credential attainment years. Our results suggest that in such countries, a policy to extend the final year of 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. In a country where the average number of years spent in school is lower and grows less quickly than in China, the benefits of such a policy could well exceed the costs.

Finally, we believe 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.

[27] These countries had more than twice the per-capita GDP than those without bunching. It is likely that productive characteristics are less observable in non-agricultural occupations (e.g. brawn as opposed to brains), and the greater signaling among richer countries we observe is consistent with a higher importance of signaling in economies with more developed non-agricultural labor markets.

References

Almond, D., Li, H., and Zhang, S. (2013). Land reform and sex selection in china. Technical report, National Bureau of Economic Research.

Angrist, J. and Krueger, A. (1991). Does compulsory school attendance affect schooling and earnings? *Quarterly Journal of Economics*, 106(4):979–1014.

Becker, G. S. (1975). *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. National Bureau of Economic Research, Inc.

Becker, G. S. and Tomes, N. (1976). Child endowments and the quantity and quality of children. *Journal of Political Economy*, 84(4):S143–S162.

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

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

Bobonis, G. J. (2009). Is the allocation of resources within the household efficient? new evidence from a randomized experiment. *Journal of Political Economy*, 117(3):453–503.

Butrymowicz, S. (2012). In china, private colleges, universities multiply to meet higher-education demand. *Washington Post*.

Cai, F., Park, A., and Zhao, Y. (2008). The chinese labor market in the reform era. In *China's Great Economic Transformation*. Cambridge University Press.

Chetty, R., Friedman, J. N., Hilger, N., Saez, E., Schanzenbach, D. W., and Yagan, D. (2011). How does your kindergarten classroom affect your earnings? evidence from project star. *The Quarterly Journal of Economics*, 126(4):1593–1660.

Chyi, H. and Zhou, B. (2014). The effects of tuition reforms on school enrollment in rural china. *Economics of Education Review*, 38:104–123.

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

Clark, D. and Royer, H. (2013). The effect of education on adult mortality and health: Evidence from britain. *American Economic Review*, 103(6):2087–2120.

Connelly, R. and Zheng, Z. (2003). Determinants of school enrollment and completion of 10 to 18 year olds in china. *Economics of Education Review*, 22(4):379–388.

Dizon-Ross, R. (2014). Parents' perceptions and children's education: Experimental evidence from malawi. Technical report, Working Paper.

Dobbie, W. and Fryer, R. G. (2013). Getting beneath the veil of effective schools: Evidence from new york city. *American Economic Journal: Applied Economics*, 5(4):28–60.

Duflo, E. (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, H., Eggleston, K. N., Rizzo, J. A., Rozelle, S., and Zeckhauser, R. J. (2012). The returns to education in china: Evidence from the 1986 compulsory education law. Working Paper 18189, National Bureau of Economic Research.

Fukunaga, K. and Hostetler, L. (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, A. and Imbens, G. (2014). Why high-order polynomials should not be used in regression discontinuity designs. Technical report, National Bureau of Economic Research.

Grenet, J. (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, E. (1999). Political change and the urban-rural gap in basic education in china, 1949-1990. *Comparative Education Review*, 43(2):193–211.

Hannum, E., Zhang, Y., and Wang, M. (2013). Why are returns to education higher for women than for men in urban china? *The China Quarterly*, 215:616–640.

Hannum, E. C., Behrman, J., Wang, M., and Liu, J. (2008). Education in the reform era. In *China's great economic transformation*. Cambridge University Press.

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

Hungerford, T. and Solon, G. (1987). Sheepskin effects in the returns to education. *Review of Economics and Statistics*, pages 175–177.

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

Jayachandran, S. (2014). Incentives to teach badly: After-school tutoring in developing countries. *Journal of Development Economics*, 108:190–205.

Karabarbounis, L. and Neiman, B. (2014). The global decline of the labor share. *The Quarterly Journal of Economics*, 129(1):61–103.

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

Lang, K. and Kropp, D. (1986). Human capital versus sorting: the effects of compulsory attendance laws. *Quarterly Journal of Economics*, pages 609–624.

Larmer, B. (2014). Inside a chinese test-prep factory. *New York Times*.

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

Lee, D. S. and Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of Economic Literature*, 48(2):281–355.

Lee, Y. S. (2011). School districting and the origins of residential land price inequality. Working paper.

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

Li, H., Li, L., Wu, B., and Xiong, Y. (2012a). The end of cheap chinese labor. *Journal of Economic Perspectives*, 26(4):57–74.

Li, H., Liu, P. W., and Zhang, J. (2012b). Estimating returns to education using twins in urban china. *Journal of Development Economics*, 97(2):494–504.

Liu, C., Zhang, L., Luo, R., Rozelle, S., Sharbono, B., and Shi, Y. (2009). Development challenges, tuition barriers, and high school education in china. *Asia Pacific Journal of Education*, 29(4):503–520.

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

Loyalka, P., Liu, C., Song, Y., Yi, H., Huang, X., Wei, J., Zhang, L., Shi, Y., Chu, J., and Rozelle, S. (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.

Lv, P. and Xie, Y. (2012). Sampling design of the chinese family panel studies. Working paper, Institute of Social Sciences Surveys, Peking University.

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

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

Meghir, C., Palme, M., and Simeonova, E. (2013). Education, cognition and health: Evidence from a social experiment. Working Paper 19002, National Bureau of Economic Research.

Ministry of Education of the People's Republic of China, B. (2013). China education statistical yearbooks.

Mo, D., Zhang, L., Yi, H., Luo, R., Rozelle, S., and Brinton, C. (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.

Mobius, M. M., Niederle, M., Niehaus, P., and Rosenblat, T. S. (2011). Managing self-confidence: Theory and experimental evidence. Technical report, National Bureau of Economic Research.

Munshi, K. and Rosenzweig, M. (2013). Networks, commitment, and competence: Caste in Indian local politics. Technical report, National Bureau of Economic Research.

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

Oreopoulos, P. (2006). Estimating average and local average treatment effects of education when compulsory schooling laws really matter. *The American Economic Review*, pages 152–175.

Pischke, J.-S. and von Wachter, T. (2008). Zero returns to compulsory schooling in germany: Evidence and interpretation. *Review of Economics and Statistics*, 90(3):592–598.

Psacharopoulos, G. and Patrinos, H. A. (2004). Returns to investment in education: a further update. *Education Economics*, 12(2):111–134.

Shi, X. (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, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3):355–374.

Stephens, M. and Yang, D.-Y. (2014). Compulsory education and the benefits of schooling. *The American Economic Review*, 104(6):1777–1792.

Tyler, J. H., Murnane, R. J., and Willett, J. B. (2000). Estimating the labor market signaling value of the ged. *Quarterly Journal of Economics*, pages 431–468.

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

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

Weiss, A. (1995). Human capital vs. signalling explanations of wages. *The Journal of Economic Perspectives*, pages 133–154.

World Bank, P. (2012). *World Development Indicators 2012*. World Bank Publications.

Zhang, J., Zhao, Y., Park, A., and Song, X. (2005). Economic returns to schooling in urban china, 1988 to 2001. *Journal of Comparative Economics*, 33(4):730–752.

Zhang, Y., Hannum, E., and Wang, M. (2008). Gender-based employment and income differences in urban china: Considering the contributions of marriage and parenthood. *Social Forces*, 86(4):1529–1560.

Figures and tables

Figure 1: Length of each credential, pre- and post-policy

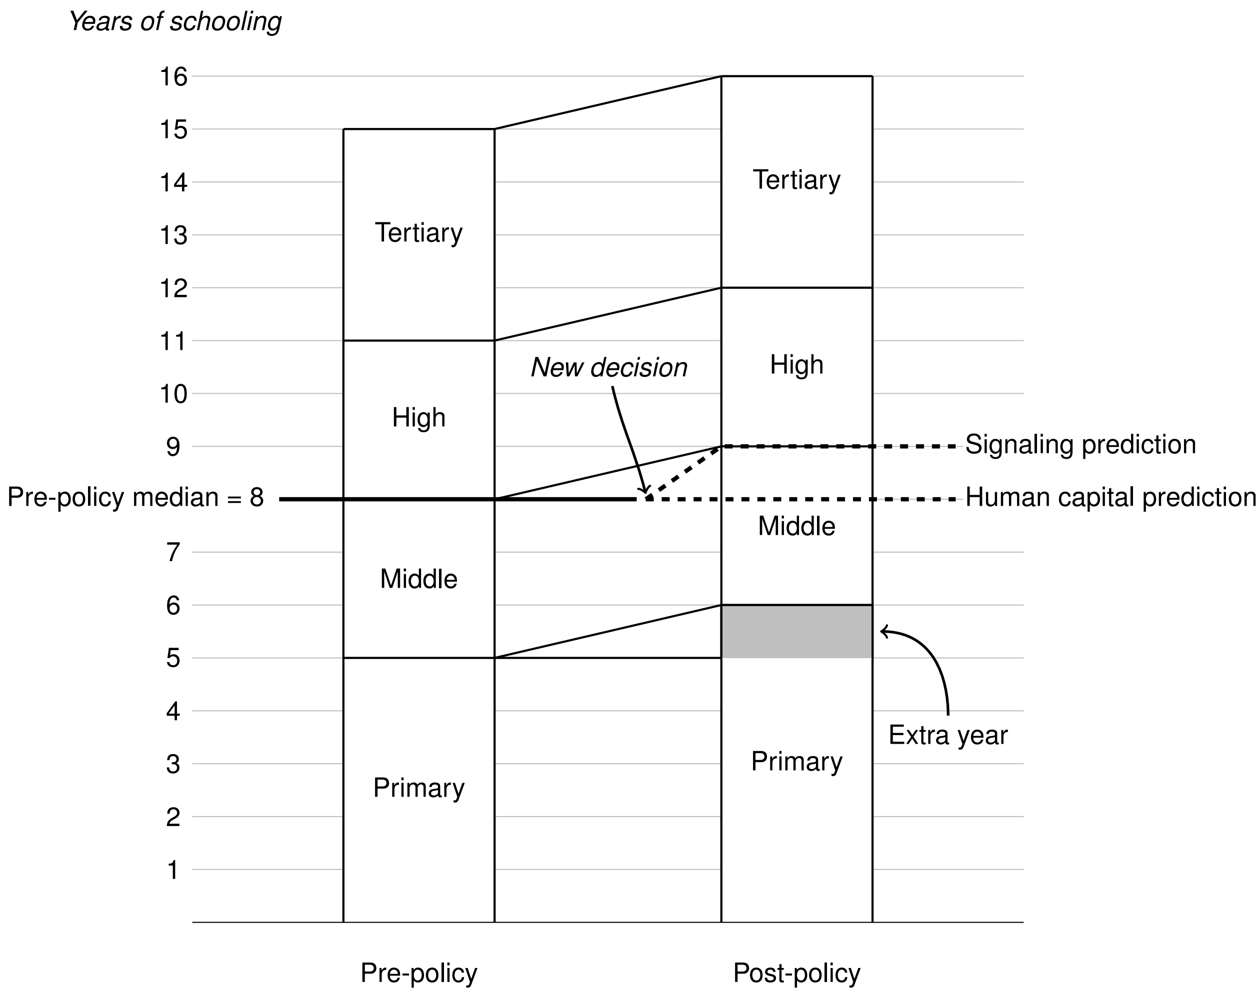

Figure 2: Timing of implementation by prefecture

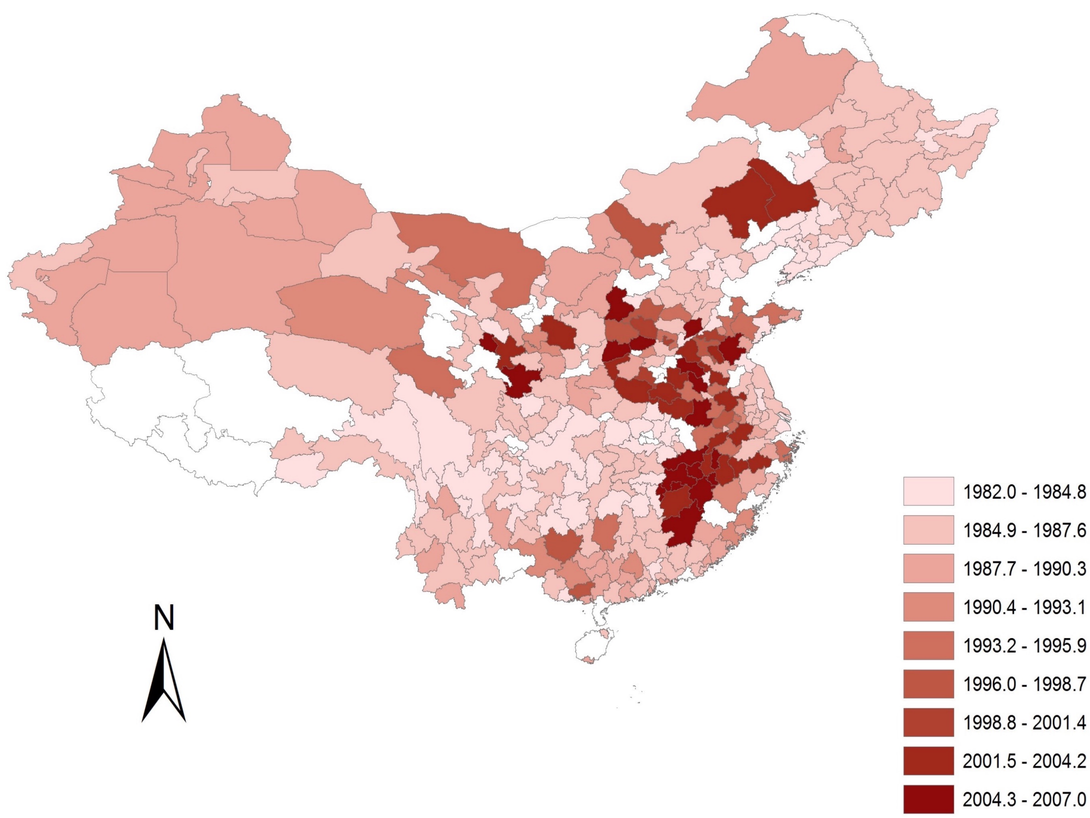

Note: In some cases, prefecture boundaries have changed since the archive was published. In these cases, we take the prefecture level-average of the treatment year in all previous prefecture capitals falling within the new prefecture.

Figure 3: Predetermined characteristics across the treatment threshold

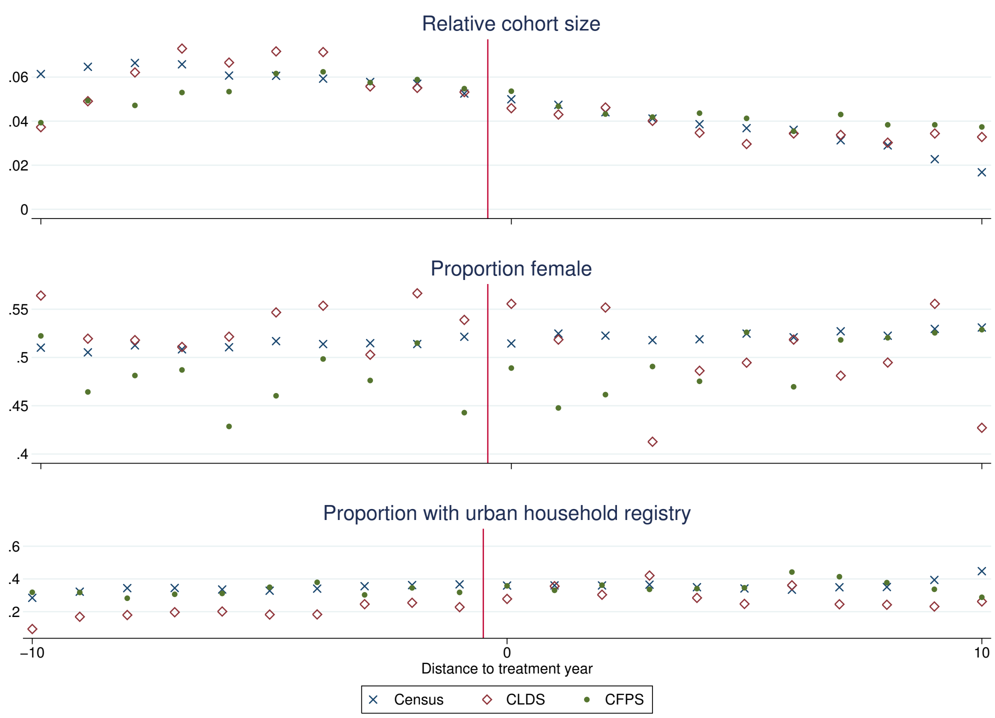

Figure 4: Proportion of population affected by policy

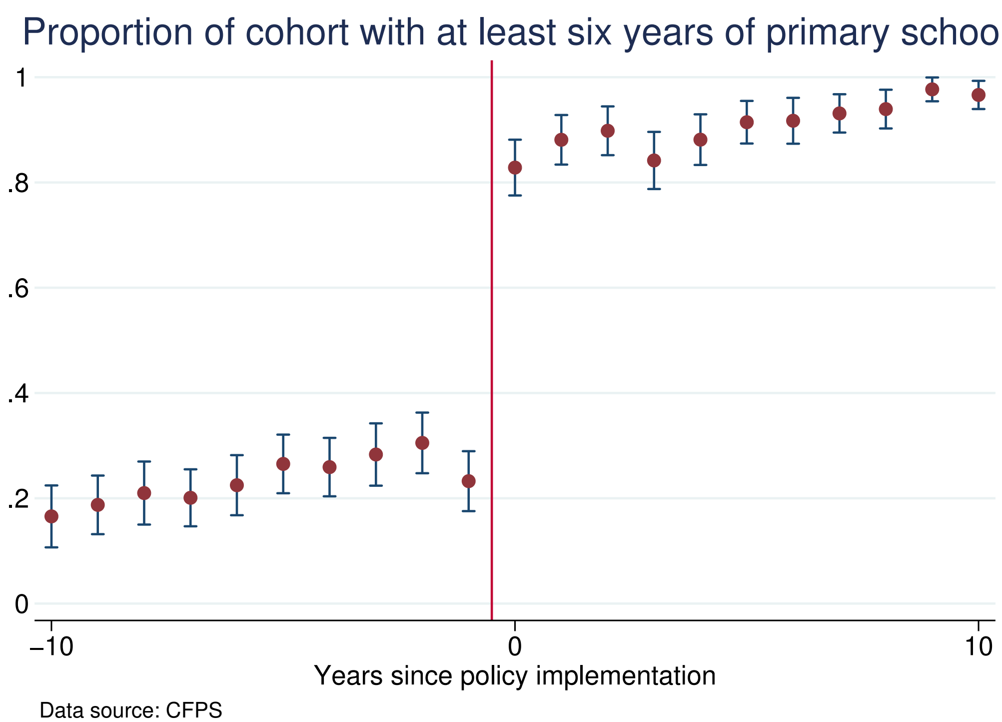

Figure 5: Distribution of post-primary education by treatment

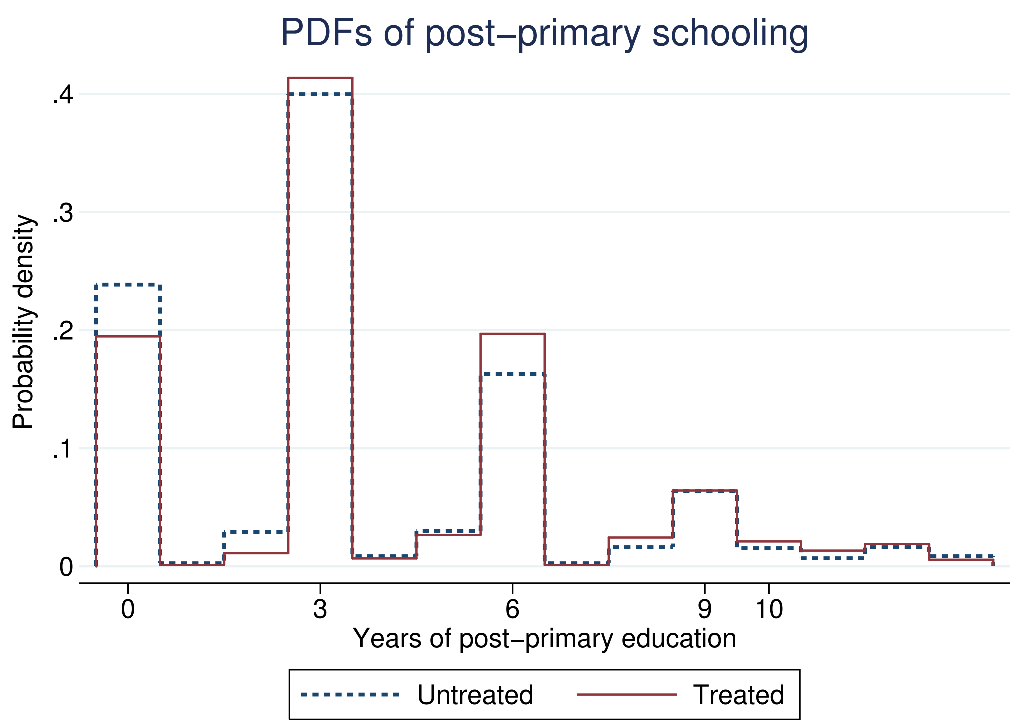

These graphs include only CFPS data from non-migrants (in counties with clarity of policy year identification = 1 or 2) graduating from primary school within five years of the policy implementation year.

Figure 6: Stability of main regression estimates across bandwidth specifications

Panel A - post-primary schooling estimates and confidence intervals

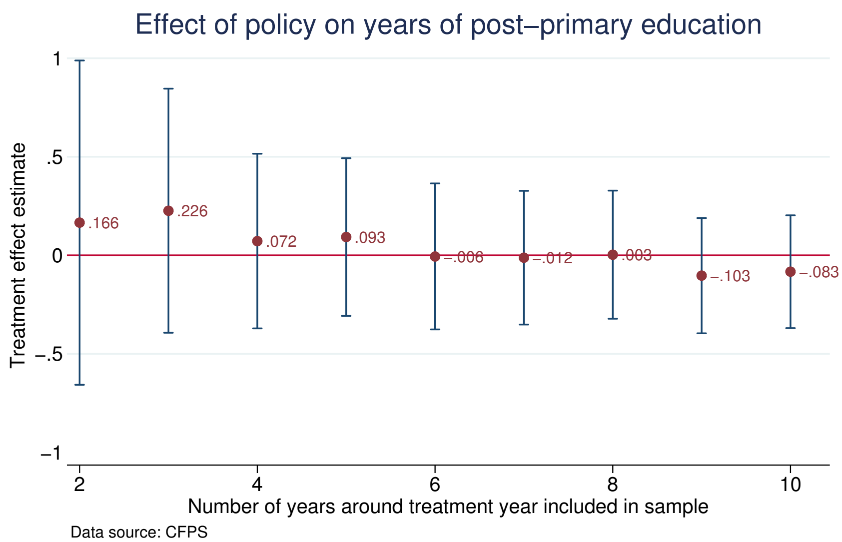

Panel B - log monthly income estimates and confidence intervals

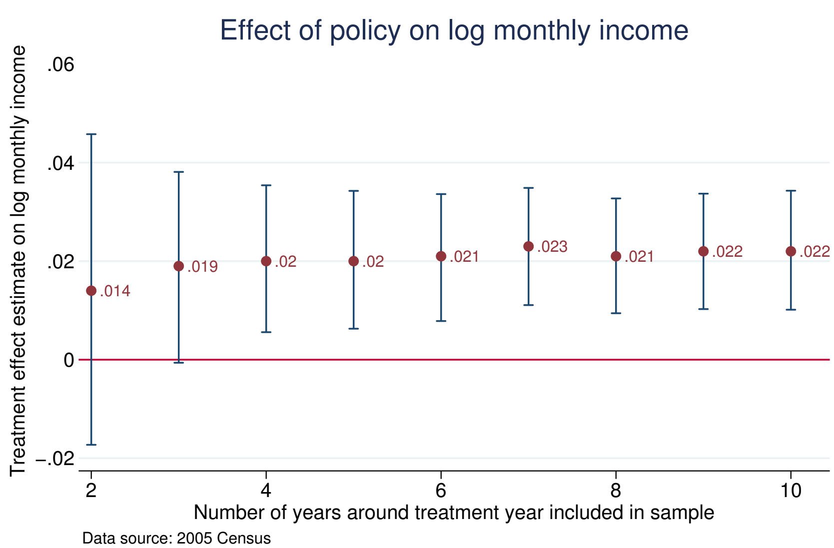

Figure 7: Estimates from 1,000 draws of placebo-policy years

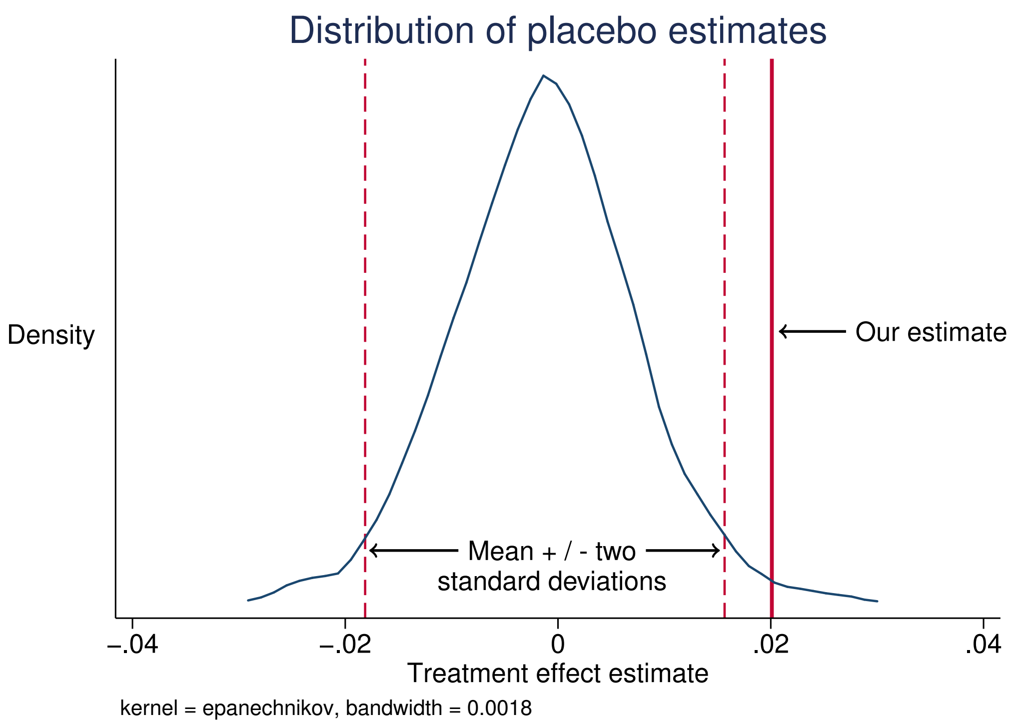

Figure 8: Kernel density of cognitive ability score, by treatment status

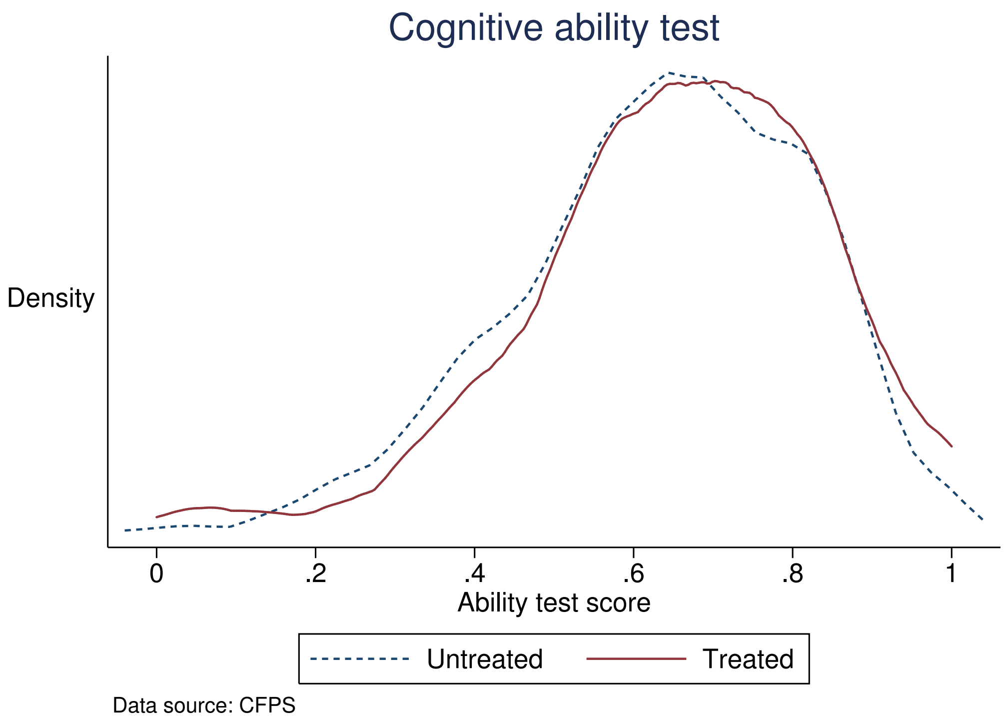

Table 1: Data sources

Data source	Year(s) collected	Sample size	Relevant contents	Policy year identification method	Geographic specificity
Census 1% sample	2005	2,585,481	Education: degree attainment Labor market: income, employment	Archival records	Prefecture*
CFPS	2010	33,600	Education: fine-grain Labor market: income, employment	Algorithm	County
CLDS	2012	16,253	Education: fine-grain Labor market: income, employment	Algorithm and archival records	Prefecture

*Note: There are multiple counties in each prefecture

Table 2: Education regressions - main sample

Outcome	CFPS	2005 Census
Probability: at least six years of primary school	0.547*** (0.029)	
Years spent in all levels of school	0.660*** (0.209)	
Years of post-primary schooling	0.093 (0.204)	
Highest credential at least middle school	0.004 (0.032)	-0.0049 (0.0030)
Highest credential at least high school	0.000 (0.026)	0.0063* (0.0033)
Dropped out of school, any level of school		-0.0002 (0.0014)

Robust standard errors in parentheses, clustered at the county/prefecture level. * p<0.10, ** p<0.05, *** p<0.01

Table 3: Education regressions - vulnerable subgroups

Outcome	Female	Rural	Rural female
Years of post-primary schooling			
CFPS	-0.2992	0.2727	-0.3409
	(0.3714)	(0.3319)	(0.5381)
Graduated from middle school			
CFPS	-0.0351	0.0589	-0.0177
	(0.0576)	(0.0618)	(0.0794)
Census	-0.0057	-0.0049	-0.0044
	(0.0043)	(0.0046)	(0.0063)
Graduated from high school			
CFPS	-0.0466	0.0131	-0.0241
	(0.0469)	(0.0202)	(0.0316)
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	946	927	412
Observations in Census sample	126,081	157,308	81,490

Robust standard errors in parentheses, clustered at the county/prefecture level. * p<0.10, ** p<0.05, *** p<0.01

Table 4: Labor market results - entire sample

Outcome	Estimates
Panel A: average treatment effects	
Currently employed	0.0026
	(0.0052)
Works for government	0.0035
	(0.0062)
Log lof 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 qualification	
Primary school	0.0657**
	(0.0312)
Middle school	0.0475***
	(0.0101)
High school	0.0311***
	(0.0096)
Tertiary schooling	-0.0066
	(0.0095)
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)

For first row, N=86,240. For all other regressions, N=66,425. Robust standard errors in parentheses, clustered at the prefecture level.
* p<0.10, ** p<0.05, *** p<0.01

Table 5: Putting our estimates in context

Credential	(1) Per-year return	(2) Credential premium	(3) Percent employed	(4) Percent of estimation sample
Primary school	-	-	62.3	3.8
Middle school	7.93	23.7	66.9	34.2
High school	8.42	25.3	77.7	32.3
Technical college	8.58	25.7	92.1	19.5
University	13.9	55.7	95.8	10.1
Our estimate of the return to a sixth year of primary school	2.03	-	-	-
Li et al. (JDE 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. Degree premia are measured relative to next highest credential: middle school relative to primary earners, high school relative to middle, and technical college and university both relative to high school.

Table 6: Cost-benefit calculation

	Assumption about heterogeneity between credential holders	
Assumption about treatment effect for rural areas	Using average treatment effect for all	Using credential-specific treatment effects
Using estimates from urban areas	-29,887	28,609
Using rural effect estimates from our data	-67,533	-173,951
Costs	146,309	-

*Estimates in millions of 2005 US Dollars.

Appendices

A1 - Figures

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

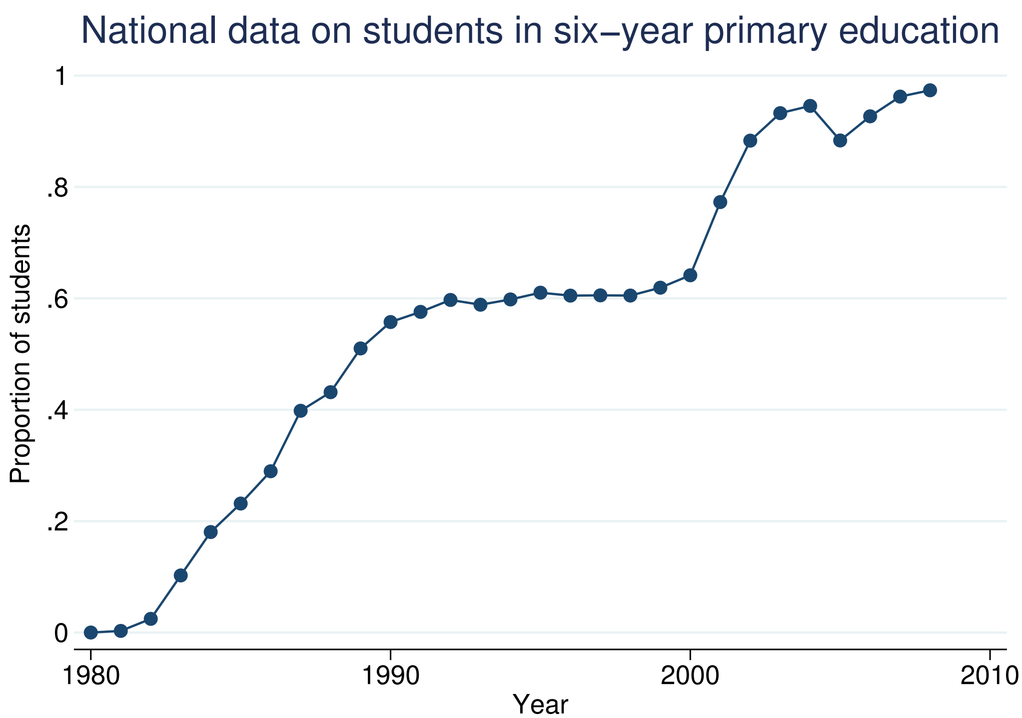

Figure A.2: National statistics and province-level CDFs of policy implementation

First 12 provinces

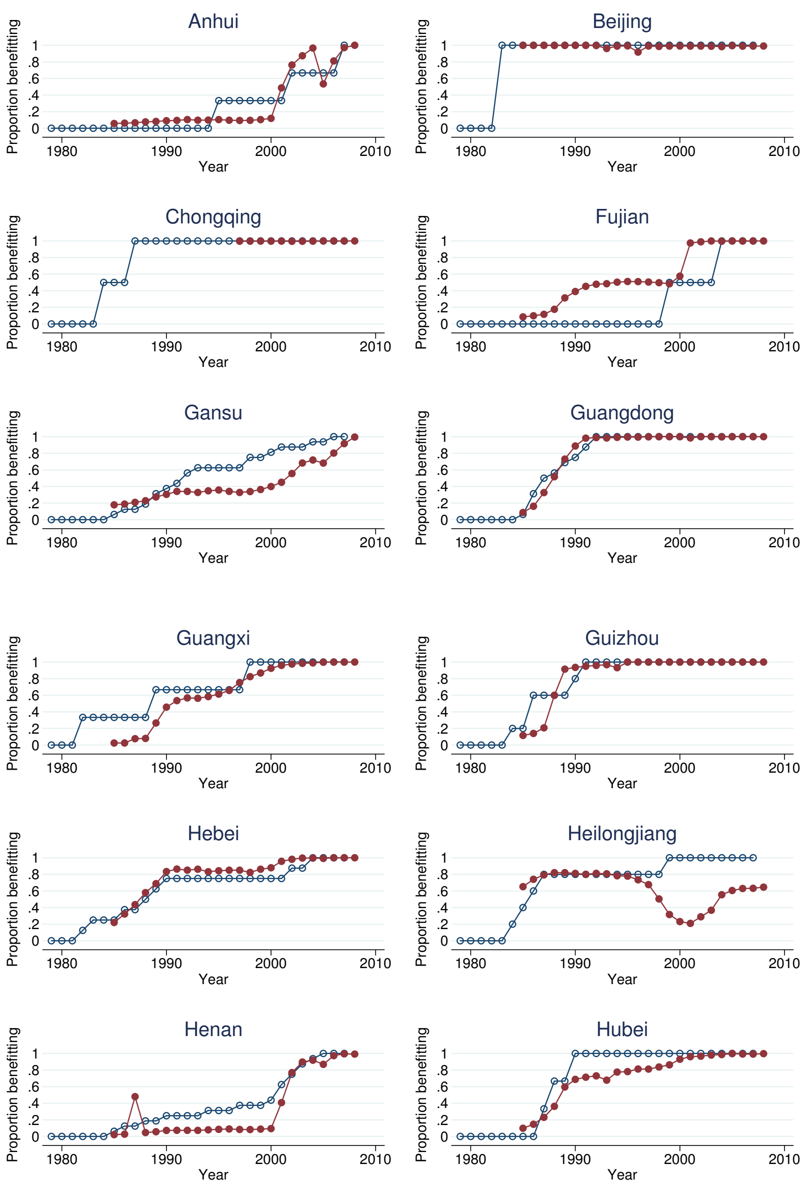

Note: hollow circles represent the cdf of proportion of treated counties in a given province by year according to county-specific implementation years identified in the CFPS data. Solid circles are from publicly available, Chinese Ministry of Education statistics on the number of schoolchildren benefiting from six year education in a given province by year. Measurement error and policy experimentation cause these to trend downward in a few cases.

Figure A.3: National statistics and province-level CDFs of policy implementation

Second 12 provinces

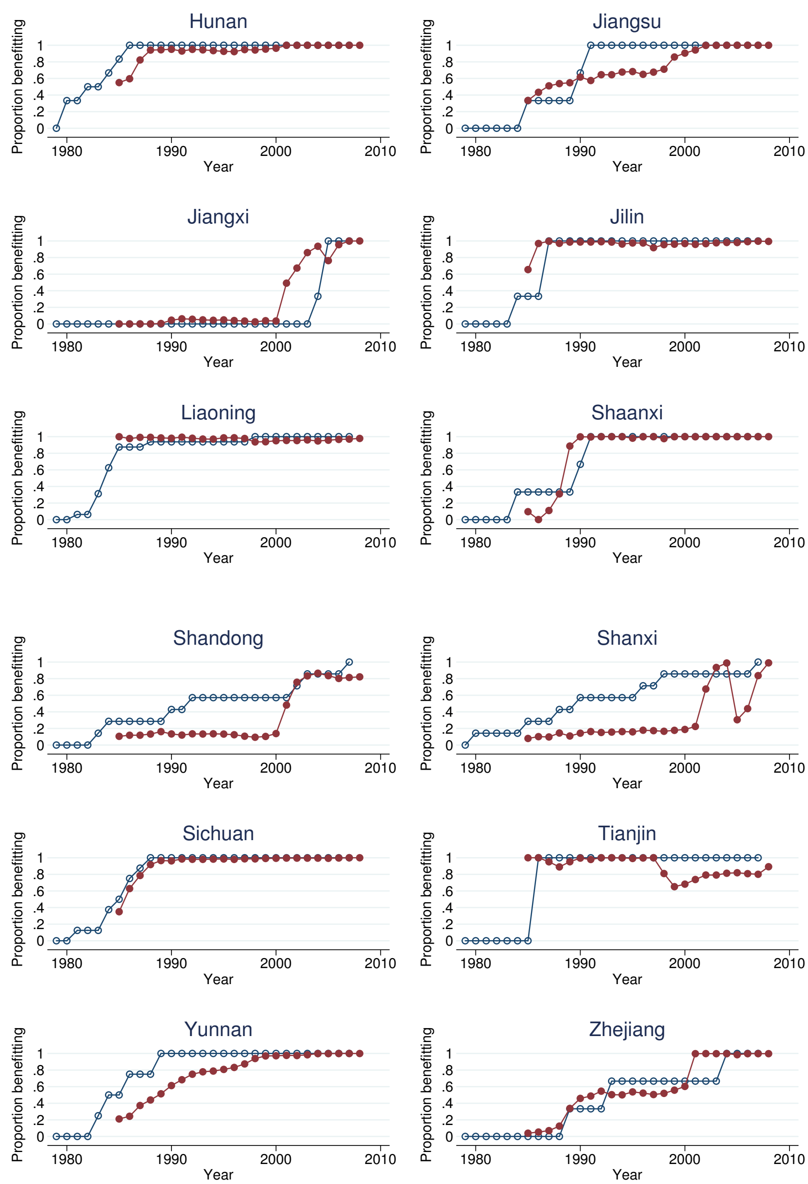

Note: hollow circles represent the cdf of proportion of treated counties in a given province by year according to county-specific implementation years identified in the CFPS data. Solid circles are from publicly available, Chinese Ministry of Education statistics on the number of schoolchildren benefiting from six year education in a given province by year. Measurement error and policy experimentation cause these to trend downward in a few cases.

Figure A.4: Example of a page from a county gazetteer

第三章 小学教育 65

第二节 学制 课程

一、学 制

威信从光绪三十三年（1907）设置小学堂以来，均执行各个时期中央政府或省的有关学制。先后实行过五、四分段的九年学制，四、三分段的七年学制，五年一贯制和四、二分段的六年学制。

1902年，清政府拟订《钦定学堂章程》，但未及实行，次年以日本学制为蓝本，重新拟订《奏定学堂章程》，于1904年1月正式公布实行，这一年为癸卯年，故又称“癸卯学制”。《癸卯学制》规定初等小学堂学制五年，高等小学堂学制四年，小学阶段共九年，儿童7周岁入学。威信1907年举办的两所初等小学堂都奉行《癸卯学制》。

民国元年（1912），南京临时政府明令废除《癸卯学制》，于次年颁布了《壬子癸丑学制》，通令全国执行（临时政府教育部于1912年制定了一个《壬子学制》，1913年又作了补充和修改，因1912年为旧历壬子年，1913年为旧历癸丑年，故称《壬子癸丑学制》）。该学制对各类学校的目的、任务、修业年限、课程设置、入学条件都作了规定。初等教育方面，规定儿童入学年龄为六周岁，修业年限则缩短为七年，初小四年，为义务教育，毕业后可升入高等小学校或乙种实业学校；高小三年，毕业后可升入中学或师范学校、甲种实业学校。威信于民国3年成立劝学所，实施新教育，兴办的初、高两级小学校也都按此学制施行。这一学制实行时间长达十年之久。

民国11年（1922），北洋政府颁布《壬戌学制》（1922年北洋政府召开学制会议，公布了《学校系统改革草案》和《新学制课程标准纲要》，对各类学校的修业年限、课程设置都作了规定。这一年为旧历壬戌年，故称《壬戌学制》），规定小学校前期四年、后期二年，学制缩短为六年，不分初等小学、高等小学，统称小学校，但初等小学仍得单独设立，为义务教育，儿童入学年龄为6足岁。威信从1923年开始实行此学制。

1950年，威信解放后，按云南省人民政府颁布的《云南省小学教育暂行实施办法（草案）》的规定执行。小学仍暂定为四、二制，初级小学招收6～9岁的学龄儿童，农村年龄稍大的也可招收入学。初小四年可单独设立，高小二年与初小合设的称完全小学。春、秋两季均可招收学生。

1952年，政务院公布的《关于改革学制的决定》和教育部颁布的《小学暂行规程》指出：目前执行的六年学制，广大劳动人民子女难于受到完全的初等教育，决定缩短修业年限，试行小学五年一贯制。当年，威信从秋季新生开始实行五年一贯制，并改为秋季始业，春季不再招生。1953年11月，政务院在《整顿和改进小学教育的指示》中规定：“关于小学五年一贯制，从执行情况看，由于师资、教材、条件准备不足，不宜继续推行，从本学期起一律停止推行。”全县小学仍沿用四、二学制。

1960年，再次进行学制改革试验，昭通专署决定威信县的新城小学、大河小学为重点试验五年一贯制学校，县又决定增加麟凤小学。秋季，三校在一年级开始进行五年

Figure A.5: Example of mean shift algorithm identifying year of policy implementation

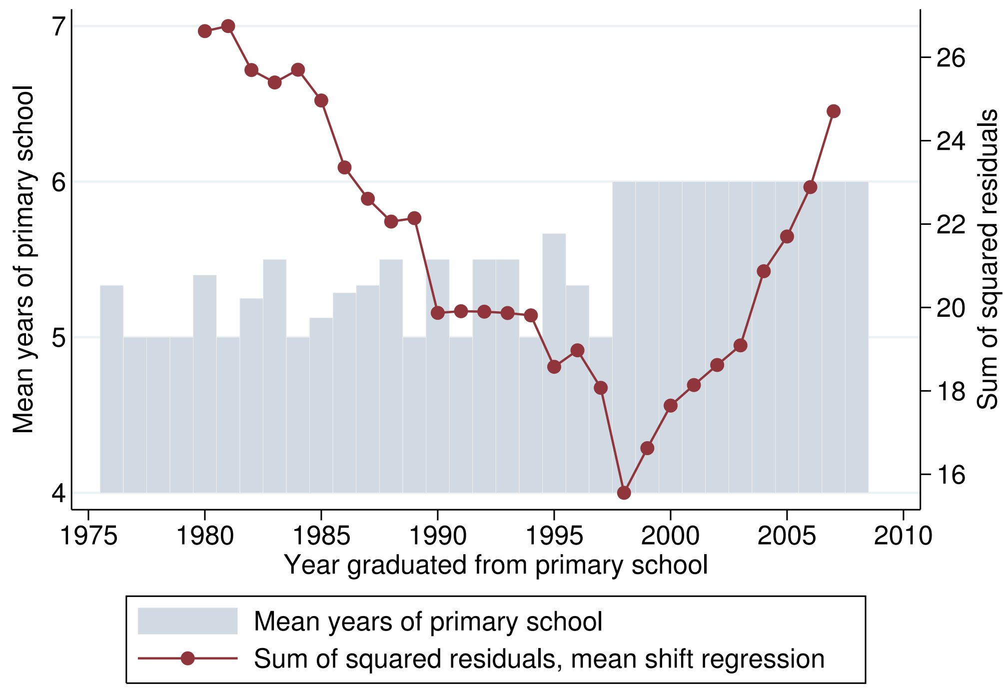

Data source: CFPS

Figure A.6: Comparing archival and algorithmically identified treatment years

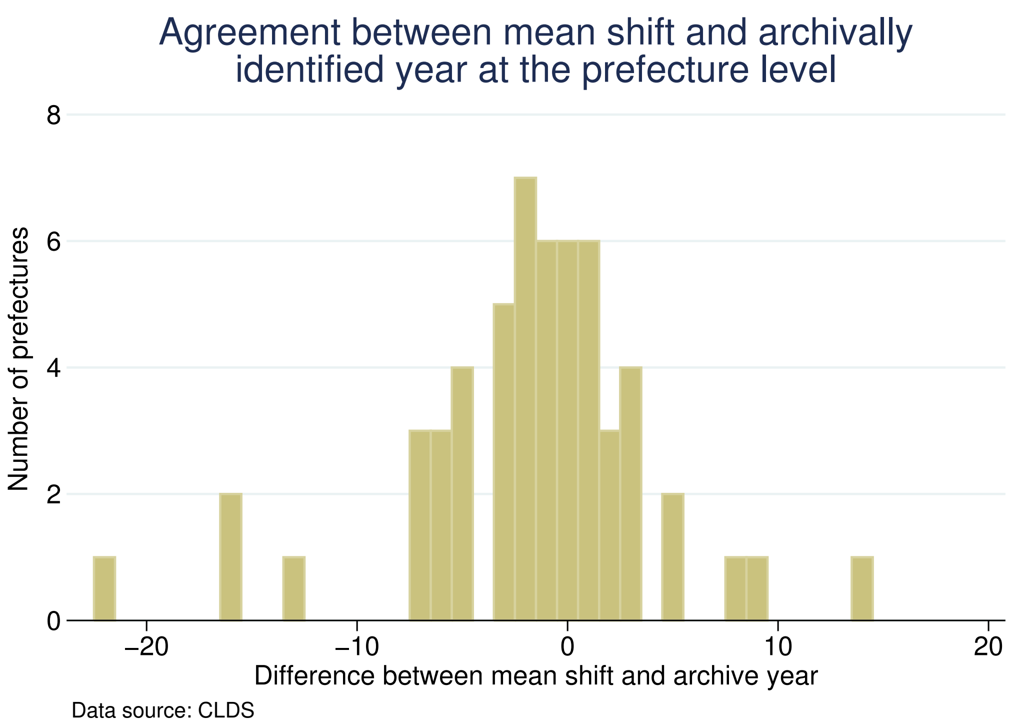

Figure A.7: Educational attainment before and after policy implementation

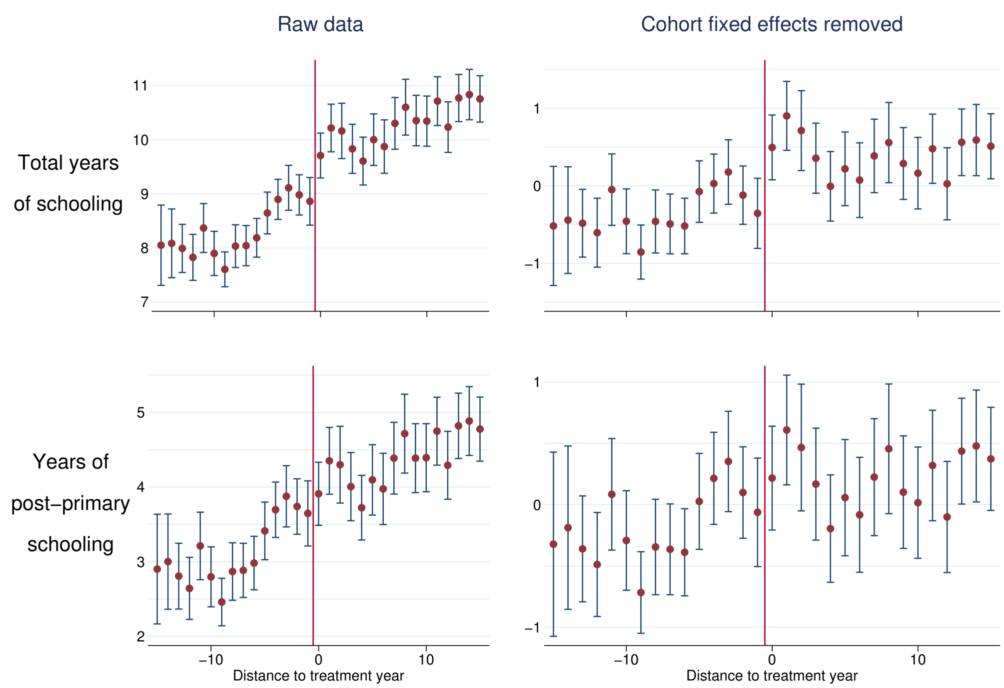

Data source: CFPS

Figure A.8: Evidence of students' ability to adjust on lower and upper margins

Panel A: Ability to adjust on lower margin

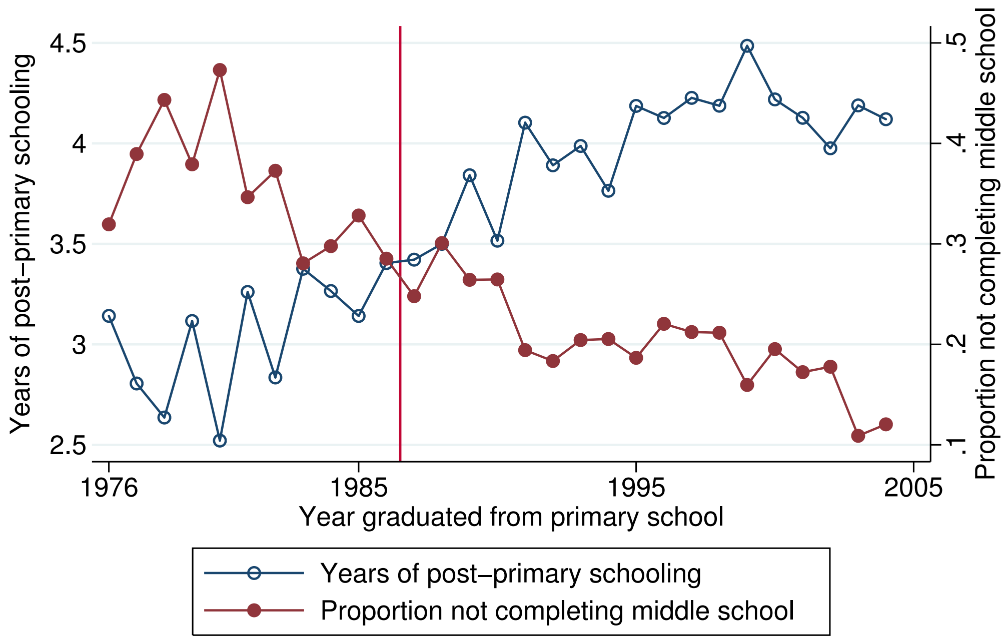

*Vertical line at annoucement of compulsory middle school

Panel B: Ability to adjust on upper margin

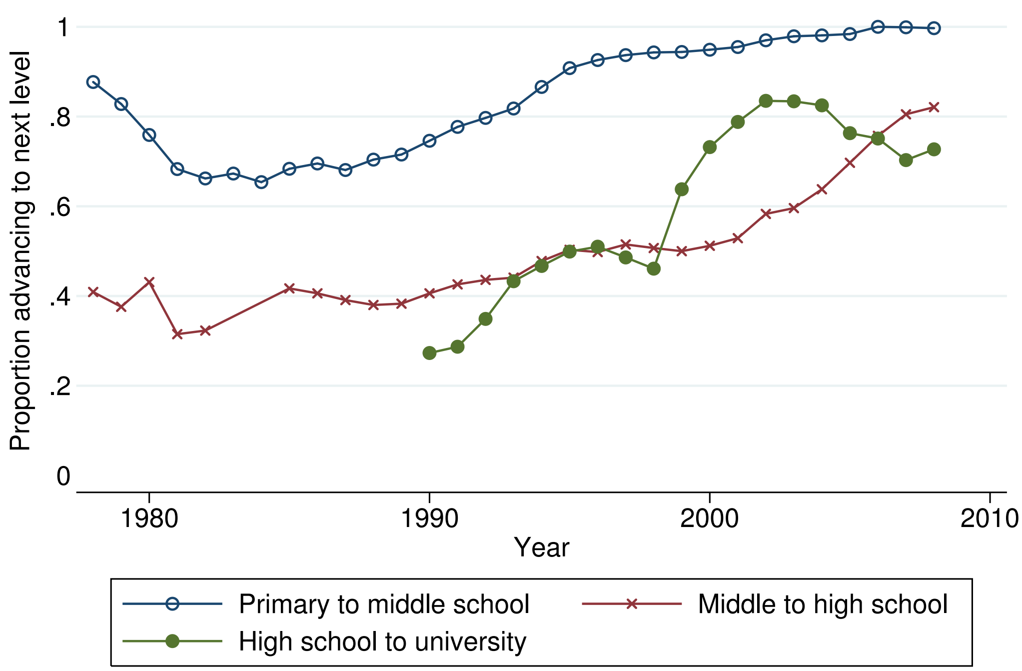

Figure A.9: PDFs of educational attainment by dominant mechanism

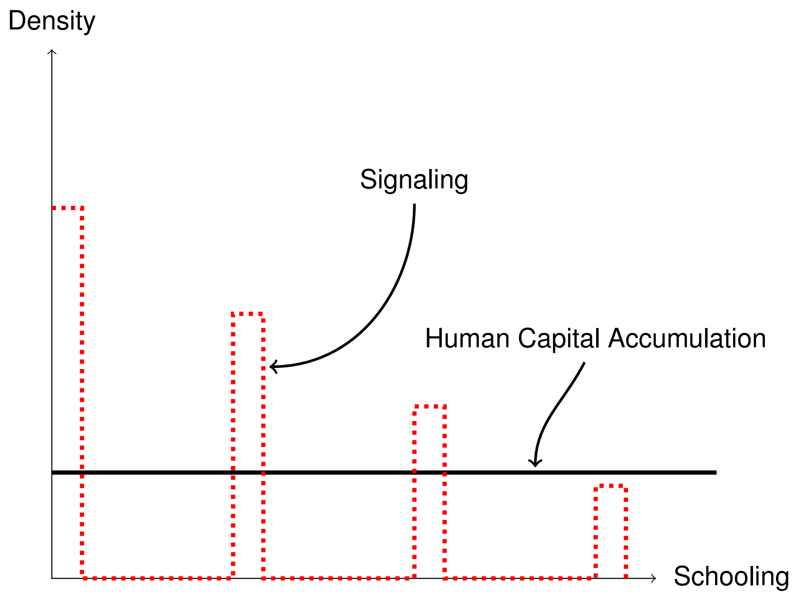

Figure A.10: Predicted policy response, by model

Human Capital Accumulation model

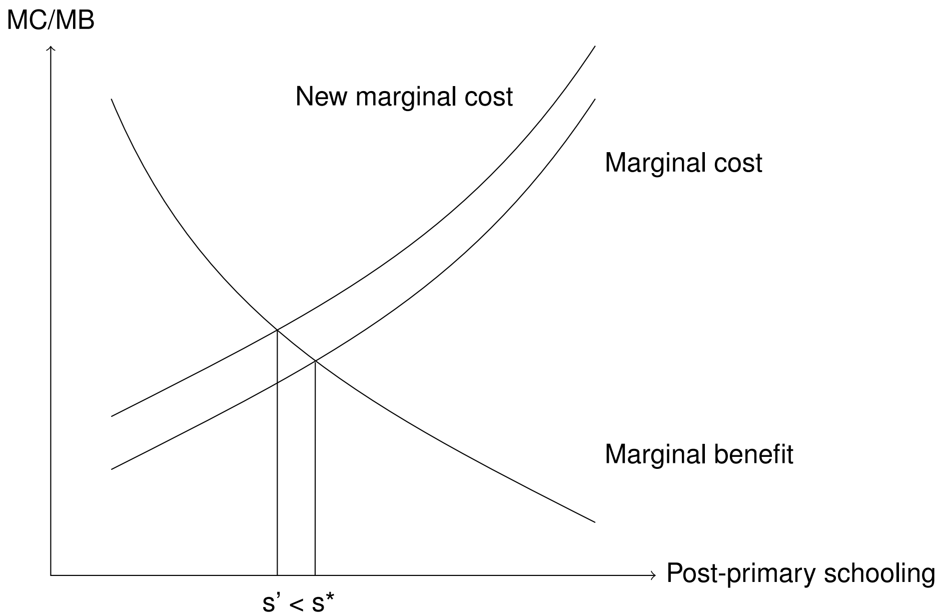

Signaling model

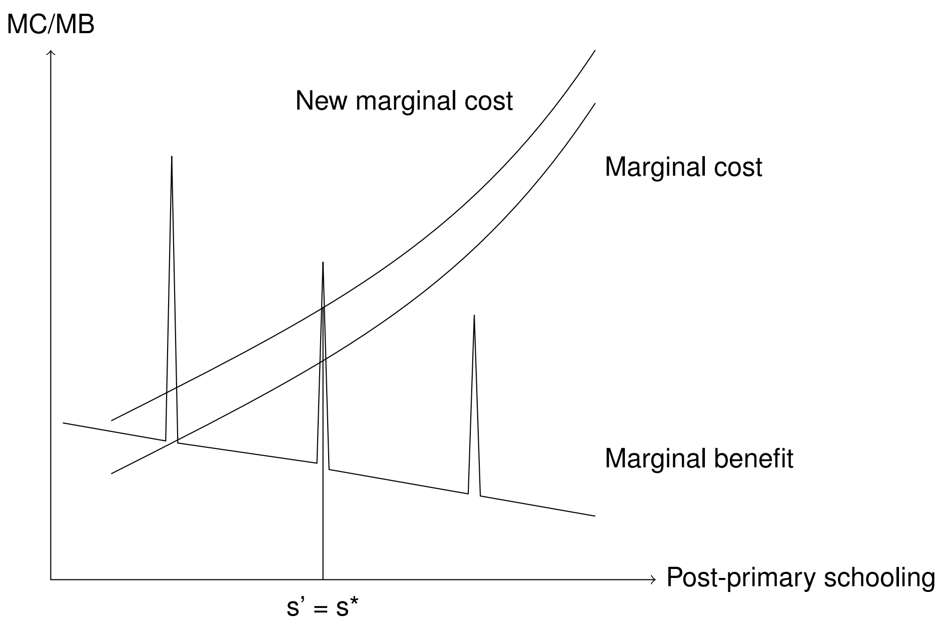

A2 - Tables

Table A.1: Anecdotes on implementation of six year primary education from across China

Region	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)

A3 Using 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 median number of years of primary school a cohort in a given county gets 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 get less primary schooling, that is, $\beta_1 < 0$, and 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 eleven 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 from us on 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 Appendix Figures A.2 and A.3. 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. Appendix Figure A.6 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-county 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.

A4 Theoretical appendix - incorporating the role of ability and the entrance exam

Here we focus on the transition from high school to college, for which the supply constraint is more likely to be binding than earlier transitions, particularly for the latter half of our study period (Connelly and Zheng, 2003). We study a continuum of families heterogeneous in two characteristics, ability endowment of the child and income. We assume income, y, and ability, a, are distributed independently and uniformly: $y \sim U[0, \overline{y}]$, $a \sim U[0, \overline{a}]$, and $a \perp y$. To simplify, in this section we assume it is always preferable to go to school if it is affordable and the student is of sufficient ability.

Here the student must have sufficient resources to ascend to the next level of schooling. There are absolute minimum levels of ability, $\underline{a}$, necessary to even potentially pass the entrance exam, and income, $\underline{y}$, necessary to pay for school should the student get in, such that if $a < \underline{a}$ or $y < \underline{y}$, then $s_t = 0$. In addition, if both are exceeded, additional expenditure is necessary for all but the brightest, the amount of which is inversely related to ability endowment[28]. This assumption captures a feature of the Chinese education system as well as of many others in East and Southeast Asia: entrance exams are high stakes and highly manipulable via test preparation (Lee, 2011; Jayachandran, 2014). This additional expenditure needed to proceed to the next level includes both money spent on tutoring, particularly exam preparation courses, and time. Time resources consist of parental time spent helping the child with her studies and time the family allows the child to spend doing homework and not housework (Zhang et al., 2008). These assumptions generate a schooling threshold in ability-income space, shown in the figure below, determined by the two absolute minima and the function determining the minimum combinations of ability and expenditure (classified as income in the proceeding figures to simplify notation) needed to pass the test.

[28]There is a large literature, starting with Becker and Tomes (1976), on whether expenditure on children's education is positively or negatively correlated with children's ability. Our assumption here is only about the likelihood of a child passing an entrance exam, and is an extension of two simple assumptions: 1) a child's likelihood of passing an entrance exam is positively related to both her ability and the amount of resources spent on exam preparation, and 2) ability and exam preparation are substitutes in the production function for passing an entrance exam.

In this framework, there are two direct effects of the compulsory primary education expansion policy, as shown in the second figure below.

EFFECT 1: $\underline{a}_{pre} > \underline{a}_{post}$. By making all children go through an extra year of primary education, the schooling reform condenses the ability distribution from the left hand side toward the right, reducing the absolute ability endowment threshold $\underline{a}$. This is supported by the results of Meghir et al. (2013), who find that a compulsory education reform in Sweden had a similar differential impact on the cognitive skills of those with initially lower ability endowments. Note that we could use a relative threshold instead of an absolute one. A relative threshold that allows only a fixed proportion of students to advance (e.g. "grading on a curve) would allow this result in the presence of noise in the entrance exam's capacity to measure ability endowments. This noise introduces a nonzero probability of failing the test inversely related to ability and money spent on tutoring.

EFFECT 2: $\bar{y}_{pre} < \bar{y}_{post}$. The policy also increases the income threshold for ascent - the extra year of schooling represents an additional year of expenditures and forgone wages to be borne for all households wanting to send their children on to further schooling. In the body of the paper we present our results testing for the second effect; we find some evidence that the most vulnerable may have been affected as predicted. As we do not have data on cognitive ability before the policy is implemented, we are unable to take the first test to the data.

Figure A.11: The schooling choice in income and ability

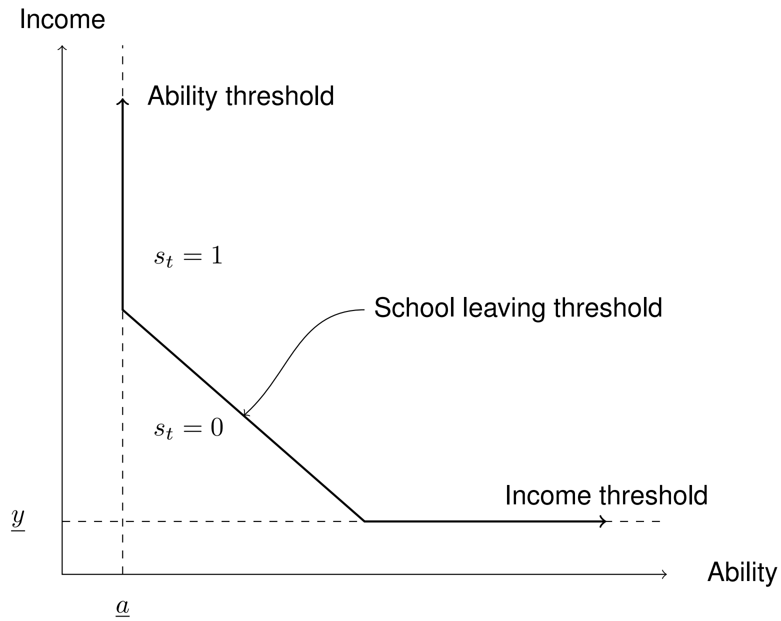

Figure A.12: Potential subgroup effects of the policy

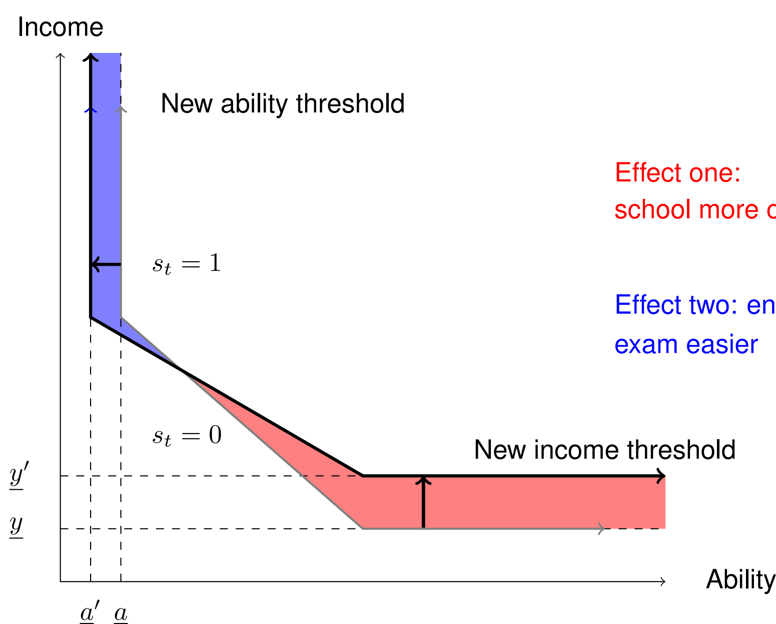

A5 Cost benefit analysis details

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 to $\frac{\alpha * GDP_t}{Workforce_t}$. 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 get 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 one of our data sets). 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 5. 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.2 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 exclude these because archival records suggest that in the short term, the current staff was used to meet most of the increased HR needs imposed by the policy and classrooms could be split, thus requiring no sizeable infrastructure additions.

In the long term, the one child policy meant that the number of children in schools was decreasing secularly, and so existing staff levels and infrastructure would be housing a smaller and smaller number of children. We assume that this would largely offset the need for long term staff and infrastructure adjustments to accomodate the policy. 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, we exclude the possible beneficial impact of the extra year of schooling on health, fertility and other non-financial outcomes.

Table A.2: 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	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	Right censor 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, QJE 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	-