



The expansion of higher education and household saving in China[☆]

Christopher Bollinger^a, Xiaozhou Ding^{b,*}, Steven Lugauer^a

^a Department of Economics, University of Kentucky, Lexington, KY 40506, USA

^b Department of Economics, Dickinson College, Carlisle, PA 17013, USA

ARTICLE INFO

JEL classification:

D12

E21

I22

J13

Keywords:

Higher education expansion

Saving

China

ABSTRACT

We examine whether access to higher education impacts household saving rates. A 2-period model of household saving decisions demonstrates why increased college opportunities induce households with children to save more. We examine this theory using survey data from Chinese households during the unprecedented education expansion. Using estimates of the change in the expected probability of college attendance, we estimate the effect on household saving rates by comparing households before and after the reform. We find that a 10-percentage point increase in the probability of going to college raises the saving rate by 5.9 percentage points.

1. Introduction

Beginning in 1999, the Chinese government rapidly expanded the previously constrained Chinese higher education system. The expansion increased the number of students allowed into Chinese universities by over 40 percent in the first year alone. While education in China is heavily subsidized, costs to families of the students were high relative to overall income and student loans were uncommon, requiring households to pay these costs out of accumulated savings or current income (Shen & Li, 2003). In recent years, education expenditures have accounted for nearly 10 percent of Chinese household expenditures (Chi & Qian, 2016). Thus, we examine evidence for an unintended consequence of the expansion: an increase in the rate of household savings. China's household saving rate has increased dramatically since economic reforms began in 1978. Before the reforms, households typically saved less than five percent of their income. Saving gradually trended up during the 1980s and 1990s, before rapidly accelerating after 2000. Now, households save over twenty-five percent of their income, on average. Several theories have been put forward to explain the high saving rates, including demographic changes (Curtis, Lugauer, & Mark, 2015; Ge, Yang, & Zhang, 2018; Imrohroglu & Zhao, 2018a; Modigliani & Cao, 2004), income uncertainty (Chamon, Liu, & Prasad, 2013), private expenditures (Chamon & Prasad, 2010),

[☆] For helpful comments and suggestions, we thank Shiyu Cheng, Chad Curtis, Donna Ginther, Hao Guo, Chenxu Hu, Olga Malkova, Jenny Minier, Min Qiang Zhao, Shunan Zhao, editor Suqin Ge, two anonymous referees, and seminar and conference participants at the University of Kentucky Macro Group Meeting, CES North America Annual Conference, Kentucky Economic Association Annual Conference, SEA Annual Meetings, WEAI Conference, SOLE Annual Meeting, Chongqing Technology and Business University, China Meeting of the Econometric Society at Jinan University, Nanjing Audit University, Biennial Conference of China Development Studies at Shanghai Jiaotong University, University of Kentucky Applied Micro Group Meeting, and International Symposium on Human Capital and Labor Markets at Central University of Finance and Economics.

* Corresponding author.

E-mail addresses: crboll@email.uky.edu (C. Bollinger), dingx@dickinson.edu (X. Ding), steven.lugauer@uky.edu (S. Lugauer).

economic reforms (He et al., 2018), and gender-related issues (Wei & Zhang, 2011; Zhou, 2014). None of the theories fully explains the changes. We provide evidence that the college expansion was an additional factor in the rise of household savings rates in China.

To motivate our empirical approach and to clearly articulate the potential connection between education expansion and saving rates, we present a simple two-period theoretical model of household saving. The model is designed to highlight college saving decisions in a two period world where there is uncertainty in college admission. The key implication is that when the probability of admission to college rises, saving rises as well, for some households. This leads us to a testable hypothesis that, conditional on other factors which would determine the baseline saving rates prior to the expansion, families experiencing a larger rise in expectation of attendance will have higher saving rates. We exploit this intuition in our empirical estimation.

We use China's 1999 higher-education expansion as a natural experiment and estimate how the policy-induced changes in the likelihood of college affected saving rates. Guided by our theoretical findings, we focus on identifying families with school aged children whose expectation of college opportunities for those children changed. These families are likely to have higher saving rates - in part to catch up - than other families, even those who would have expected to send their child to college prior to the expansion. We use 1995 and 2002 data from the China Household Income Project (CHIP). We use four distinct samples: families with college aged children before and after the expansion and families with school aged children before and after the expansion. We use the college aged samples to estimate models of college attendance and then use the primary school aged samples to estimate the savings model, using the change in probability to measure the intensity of treatment. We estimate the savings models both prior to the expansion and after the expansion. Much of our variation derives from differences in the timing and magnitude of the expansion across provinces. While the expansion has been large across all provinces, in 2002, there was substantial variation.

Our approach is consistent with other literature examining the impact of the higher education expansion. For example, Che and Zhang (2018) use an approach similar to ours to show that the Chinese college expansion increased human capital and improved productivity. Feng and Xia (2018) examine how the college expansion affected technology adoption in Chinese firms. Also, see Li and Xing (2010), Knight, Deng, and Li (2017), Li et al. (2017) and Hu and Bollinger (2017, 2021). These papers primarily focus on labor market implications, while we connect college expansion to saving behavior.

Consistent with the theoretical model, our main finding is that the expansion has an economically and statistically significant effect on household saving rates. A 10 percentage point increase in the college probability increases the average household's saving rate by 5.9 percentage points. This result provides evidence that the education expansion increased saving for households with school-age children. Our results appear to be driven by households experiencing a large increase in the possibility of college, while the few households who might always have planned on college had similar savings. Our main findings are robust to a variety of alternative specifications including interactions with important demographic variables, omitting provinces with the largest changes in enrollment rates and including households without a dependent child. We also carefully consider other policy changes that occurred in China, including housing reform, public healthcare reform, and the reorganization of State-Owned Enterprises.

The rest of our paper is organized as follows. Section 2 documents the saving rates among Chinese households and provides background on the education expansion. Section 3 presents the model of household saving that motivates our empirical analysis. Section 4 introduces the data, and Section 5 discusses our estimation strategy. Section 6 presents the main empirical findings. Section 7 contains additional analysis, and Section 8 concludes.

2. Saving and higher education in China

We are interested in how the expansion of higher education affects household saving decisions. In this section, we document the large increase in Chinese household saving rates over time and provide background on the relevant changes to education policy. The remainder of the paper links these two phenomena, both theoretically and empirically.

2.1. Household saving rates in China

Throughout the paper, we define the household saving rate as

$$\text{Saving rate} = \frac{\text{Income} - \text{Expenditure}}{\text{Income}} = 1 - \frac{\text{Expenditure}}{\text{Income}}. \quad (1)$$

Income is measured by self-reported total annual income, including wages, other income, and asset returns. Expenditure is measured by total consumption expenditure. Both variables are available in China Household Income Project surveys. We use current (unadjusted) prices throughout, except where noted. Fig. 1 plots the urban household saving rate in China from 1981 to 2011 using National Bureau of Statistics data. Saving slowly trended up before 2000; after that, though, the saving rate exploded and the slope of saving rate increased as well. In recent years, the typical Chinese household has been saving 30 percent of their income, a truly astonishing number. There is wide-spread interest in understanding why Chinese households save so much.

Several theories have been put forward. One strand of the literature has focused on the economic reforms that have been ongoing since 1978. Chamon et al. (2013) argue that the reforms have led to uncertainty over income, pensions, and healthcare, inducing Chinese households to save more for precautionary reasons.¹ As an example of this, He et al. (2018) show that the reform of

¹ See Chamon and Prasad (2010) for more on how the burden for many expenditures shifted away from collectives and onto households. Also, see Yoo and Giles (2007).




Fig. 1. Household saving rates. *Notes:* The data is from the China yearly statistical book and the national bureau of statistics. The variables used to construct the saving rate are per capita total income of urban households (y) and per capita annual cash consumption expenditure of urban households (e) at current prices. The household saving equals $(y - e)/y$.

state-owned enterprises in the 1990s increased unemployment and other risks for Chinese households, generating a dramatic increase in saving. Similarly, [Choi, Lugauer, and Mark \(2017\)](#) show that high income growth coupled with high income uncertainty helps to explain why Chinese households currently save so much more than US households. Other papers, including [Song and Yang \(2010\)](#), [Lugauer and Mark \(2013\)](#), [Song et al. \(2015\)](#), and [Curtis \(2016\)](#), have studied the interaction between saving and aggregate growth within the context of China.

A second strand of the literature has focused on the implications of China's fertility policies. [Modigliani and Cao \(2004\)](#) and [Curtis et al. \(2015\)](#) appeal to the life-cycle hypothesis of household saving to argue that the large demographic changes, caused in part by the One Child Policy, have increased saving rates over time.² Several papers have documented specific examples of how demographics impact saving in China. [Imrohoroglu and Zhao \(2018b\)](#) argue that increases in expected old-age health costs and the aging population act together to increase savings. [Rosenzweig and Zhang \(2014\)](#) present evidence showing that the increase in the prevalence of co-residing elderly parents helps to explain household saving behavior. [Wei and Zhang \(2011\)](#) and [Zhou \(2014\)](#) show that a gender imbalance in the sex-ratio has interacted with life-cycle considerations to impact saving.

Our paper is related to these previous studies, in that the cost for higher education has shifted to families and saving for college varies over the life-cycle.³ China radically and very rapidly altered its higher-education policies, which allows us to study households that experienced a sudden change in their expectations for sending their children to college. We next provide background on these policy changes.

2.2. Background on higher education enrollment policy

College admission in China is based purely on the Gaokao (college entrance exam) administered every year in June. Students are admitted on a provincial basis, and the Ministry of Education decides the number of students from each province through the coordination between local governments. The enrollment rates vary substantially across provinces. So, students compete for admission with other applicants within their province. If a high school graduate fails to get into a college, they can retake the exam and reapply in the following year. Due to the Cultural Revolution, the entrance exam was suspended between 1968 and 1978. In 1978, China initiated sweeping economic reforms, including reinstating the college entrance exam and enrolling new students into universities.

[Fig. 2](#) shows the number of students matriculating, total number enrolled, the number of college teachers and staff, and the number of higher institutions by year. There were over 400,000 new students enrolled in 1978, and about 850,000 students total. The number of students gradually increased until 1999. Then, China instituted the higher education expansion, and enrollment exploded. In 1999 alone, the number of newly enrolled students increased by more than 40 percent. Over 1.5 million high school students began college in 1999, and by 2012 nearly 7 million students were starting college each year. [Fig. 3](#) plots this unprecedented increase in college enrollment alongside the national enrollment rate. The jump in the enrollment rate after 1999 (from below 40 percent to over 60) is readily apparent. Our analysis focuses on this sudden policy change.

The central Chinese government designed and controlled the 1999 higher education expansion, and it reflected the government's political objectives. China had experienced various social and economic headwinds in the 1990s, including a sizable reduction in

² Also, see [Bairliya, Miller, and Saxena \(2018\)](#), [Banerjee et al. \(2015\)](#), [Chao, Laffargue, and Yu \(2011\)](#), [Choukhmane, Coeurdacier, and Jin \(2018\)](#), [Curtis, Lugauer, and Mark \(2017\)](#), [Jia et al. \(2021\)](#), and [Lugauer, Ni, and Yin \(2019\)](#).

³ In a paper related to ours, [Chen and Yang \(2012\)](#) use the education reform in China to test the theory of precautionary saving. Also, note that housing in China shares these characteristics; see [Chen, Yang, and Zhong \(2020\)](#).

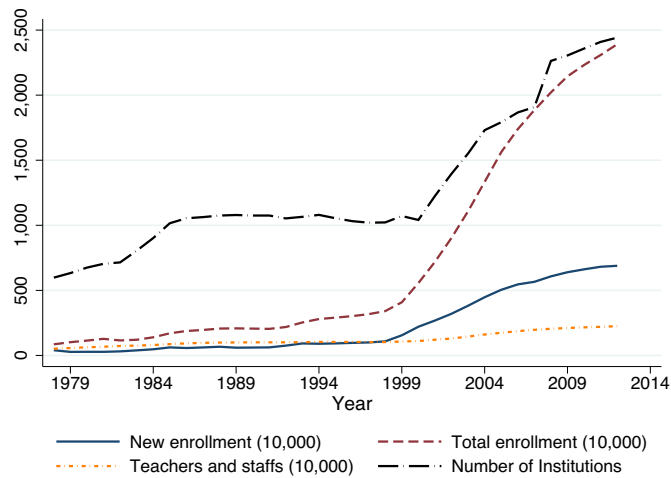


Fig. 2. China higher education expansion. *Notes:* The data comes from the China yearly statistical book and the national bureau of statistics. The variables are in units of 10,000, except the number of institutions, which is unadjusted.

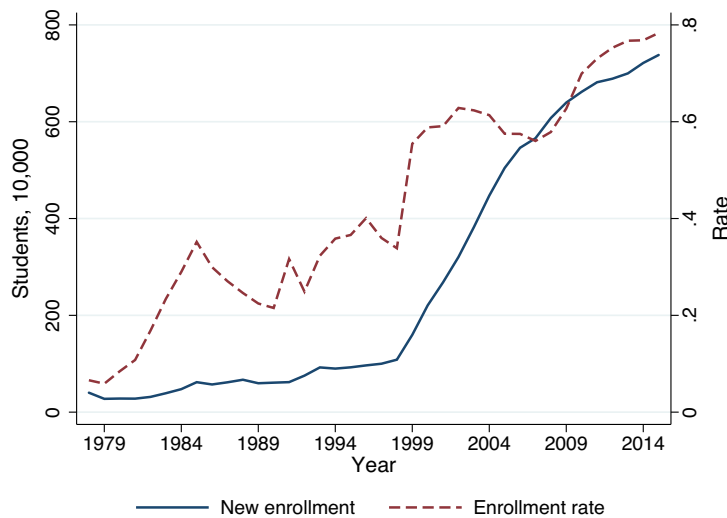


Fig. 3. New students enrolled and the enrollment rate. *Notes:* The enrollment data is from the China yearly statistical book and the national bureau of statistics. The enrollment rate is calculated by dividing the number of newly enrolled students by the number of national college entrance examination takers. The number of exam takers was collected from <http://news.koolearn.com/20180606/1152629.html>.

employment at State-Owned Enterprises. The children of the large Chinese baby-boom generation were facing a bleak employment outlook, and the government was urged to find a solution.⁴ According to lore, Min Tang, an economist at the Asian Development Bank in China, proposed enrollment expansion in a November 1998 letter to Premier Zhu Rongji. The hope was that doubling college enrollment within three years would stimulate investment in services, construction, and other related industries and would ultimately increase consumption (Wan, 2006). It was also suggested that households could use their savings to pay for college tuition and expenses. The plan was implemented almost immediately. So fast was the enrollment policy changed that typical Chinese households could not have anticipated its timing nor size. Moreover, as we show below, the policy had a differential impact across households. The increase in enrollment rates differed by province as well as by other observable household characteristics. We exploit the massive policy change to estimate how the expansion in higher education affected household saving behaviors.

⁴ Even though the One Child Policy went into effect before 1980, the sheer size of the previous generation meant that in aggregate they still had many children reaching adulthood.

The link between education and saving that we propose is a simple one. Households had to save in order to afford college tuition and expenses. Student loans were not readily available (Wang et al., 2014), and while the prevalence of college increased, so did the cost.⁵ Fig. 4 plots the total annual tuition and fees per student (in current prices) for different types of education over time. Due to compulsory education laws, the cost of attending elementary school and middle school remained close to zero. In contrast, the already high tuition and fees for college went up by a factor of four. Education was (and remains) a major expenditure category for Chinese households (Chamon & Prasad, 2010). According to Yao et al. (2011), most Chinese households save for education motives, and Wei and Zhang (2011) show that 76 percent of single-female-child households save for education.

The college expansion reform implemented by the Ministry of Education of China serves as a good natural experiment for analyzing the impact of education expansion on household saving. First and foremost it was a large exogenous shock. Second, the magnitude of the expansion varied by province even though it was implemented nation-wide. Each province received a different “quota” to determine how many students were to be admitted, and families could not easily relocate due to the Hukou registration system. Thus, the change in the expected probability of attending college depended on geography; however, as we show below, other household characteristics mattered, as well. Our empirical strategy leverages the exogenous variation in expected enrollment rates at the household level. Before detailing our approach for analyzing the data, we next present a theory for why an expansion in higher education relates to household saving behavior.

3. A model of households saving for college

This section contains a simple two-period model of household saving decisions. Our goal is to motivate our empirical analysis by presenting an explicit theory showing how enrollment rates can impact saving decisions. The key choice by households in the model is whether or not to save enough to cover college tuition and expenses. An expansion in college opportunities induces some families to increase their savings.

For tractability, we consider households with log utility over consumption C in period 1 and expected consumption C' in period 2.

$$U = \ln C + E[\ln C'] \quad (2)$$

The household receives income Y in period 1, which is split between consumption and saving S . Thus, the budget constraint in the first period is

$$C + S = Y, \quad (3)$$

where households begin life with no assets and saving must be positive. We abstract from discounting future utility and set the return on saving to zero, as these considerations do not materially alter the analysis.

The household saves in order to consume in period 2, but also to potentially pay college tuition for their child. The saving decision is made before knowing whether or not their child will be admitted into college. However, the household does know the probability p that their child gets in (where $0 < p \leq 1$). Within the model, we interpret the education expansion as an increase in p . When a child is accepted into college, the household can pay tuition and expenses τ out of savings. If the household has not saved enough, $S < \tau$, then the child cannot attend college. The household's second period budget constraint is

$$C' = S, \quad (4)$$

if $S < \tau$. If $S \geq \tau$, then, with probability $1 - p$ the budget constraint remains $C' = S$. But, with probability p , if $S \geq \tau$, then the second period budget constraint becomes

$$C' = S + \theta - \tau, \quad (5)$$

where θ is the benefit from sending a child to college.

We assume that the benefits of sending a child to college outweigh the costs; thus, $\theta > \tau > 0$, and if $S \geq \tau$, an admitted child always enrolls in college. There are many other ways to model the benefits from college, but this modeling choice is tractable and easy to interpret. The household must decide whether it wants to consume less in period 1 in order to have enough saved up to potentially pay college tuition τ . The benefit θ can be taken literally as the incremental increase in old-age support from having a college-educated child, but also more broadly to include non-monetary benefits.

With this set-up, we can examine how household saving decisions change in response to the policy of increasing the enrollment rate p . Consider a household that is saving too little to send their child to college, $S = Y/2 < \tau$. For this household, the marginal utility of consuming in period 1 exceeds the expected marginal utility of saving for college. Even if their child is accepted into college, the household will not obtain θ in the second period. However, for this low-saving household, an increase in p could induce the household to save more (enough to cover tuition τ). That is, an increase in p increases the expected utility of saving for college. The exact p in which the household is indifferent between saving τ (and possibly sending their child to college) and continuing to save less than τ is given by the following proposition.

⁵ Credit constraints often bind even when loans can be obtained by some. See Keane and Wolpin (2001), Cameron and Taber (2004), Carniero and Heckman (2009), Brown, Scholz, and Seshadri (2012), Sun and Yannelis (2016) and Malkova and Braga (2018) for example.




Fig. 4. Cost of education. *Notes:* The data comes from the China yearly statistical book and the national bureau of statistics. The reported values are per-student at current (unadjusted) prices.

Proposition 1. (Threshold p) For a household with $Y < 2\tau$, the expected utility from saving τ and saving less than τ become exactly equal at $p^t > 0$, where the threshold p^t is given by:

$$p^t = \frac{\ln \left[\frac{y^2}{4(Y-\tau)\tau} \right]}{\ln \frac{\theta}{\tau}} \tag{6}$$

The key implication from Proposition 1 is that as higher education is expanded (p increases), a low-saving household (saving less than τ) will start saving τ once p reaches the threshold p^t .⁶ This theoretical result helps to motivate our main empirical strategy, in which we estimate the idiosyncratic increases in p across households and quantify how changes in p impact saving.

An increase in p may not impact all households in the same way, though. Households with older children or no children might be affected differently, or not at all. The model could be expanded to consider these additional household types; however, the remainder of the paper focuses on households with young children. Even among households with young children, the impact could differ. Some families, experiencing a smaller increase in p , may not hit their threshold, leaving their saving behavior unchanged.⁷ Our empirical analysis exploits variation in the (estimated) changes in p across households (as suggested by the theoretical model) to estimate how the change in college probability affects household saving. We introduce the data next and then provide the full details of our empirical strategy.

4. The data

We use household level information from the China Household Income Project (CHIP).⁸ The CHIP consists of repeated cross-sections of data from household surveys that were conducted in five waves across 12 provinces. It is the most widely-used micro data set on Chinese households. The survey contains questions on income and expenditures, as well as other household characteristics such as geographic location, number of children, and educational attainment. We use the 1995 and 2002 waves of urban households, which bracket the 1999 college expansion.

The raw data contains 6929 household observations in 1995 and 6835 in 2002. We focus on households with school-age children,

⁶ We derive Proposition 1 in Section A of the Online Appendix in a straightforward way. We find the threshold p by equating the lifetime utilities from either saving for college or not. The Appendix also shows that the threshold p^t decreases with the benefits derived from college θ , increases with college costs τ , and decreases with income Y , over the relevant range of variable values.

⁷ Another set of households (e.g. with $Y > 2\tau$) may have been saving enough to send their children to college, $S \geq \tau$, prior to the education expansion. These high-income households could even reduce their saving (but only as low as τ) as p goes up. We show this in the Appendix, and the result is intuitive. An increase in p for a household already saving more than τ merely increases the chance of obtaining θ . This increase in expected income in period 2 lets the household consume more (and save less) in period 1.

⁸ As mentioned above, the education expansion likely impacted some types of households differently than those with young children. Therefore, using micro data is critical to our identification strategy; it also allows us to net out other factors affecting saving that might be correlated with the education expansion (e.g. income). We return to these issues below.

Table 1
Summary statistics by year and children's age.

	School-age		College-age	
	1995	2002	1995	2002
Saving rate	0.107 (0.226)	0.142 (0.342)	0.141 (0.238)	0.128 (0.334)
Child college attainment dummy			0.101 (0.301)	0.183 (0.387)
<i>Controls</i>				
Number of kids	1.111 (0.323)	1.076 (0.267)	1.631 (0.630)	1.159 (0.404)
Number of elderly	0.0838 (0.313)	0.0878 (0.340)	0.0755 (0.279)	0.0647 (0.262)
Private house	0.410 (0.492)	0.782 (0.413)	0.475 (0.500)	0.757 (0.429)
Housing accumulation fund	0.441 (0.497)	0.543 (0.498)	0.424 (0.494)	0.540 (0.499)
Age	39.74 (5.058)	39.94 (4.911)	50.10 (4.638)	47.27 (3.273)
Gender	0.633 (0.482)	0.652 (0.477)	0.664 (0.472)	0.638 (0.481)
Currently employed	0.980 (0.140)	0.923 (0.267)	0.860 (0.347)	0.853 (0.354)
College degree	0.066 (0.247)	0.110 (0.313)	0.100 (0.300)	0.050 (0.219)
SOE job	0.822 (0.383)	0.325 (0.468)	0.825 (0.380)	0.342 (0.475)
Public health	0.703 (0.457)	0.636 (0.481)	0.709 (0.454)	0.680 (0.467)
Employment tenure	15.80 (7.425)	13.58 (8.613)	21.41 (10.23)	16.32 (11.60)
Years of schooling	10.45 (2.984)	11.46 (3.004)	10.32 (3.651)	10.29 (2.881)
Spouse years of schooling	9.657 (3.476)	10.67 (3.571)	8.714 (4.148)	9.582 (3.439)
Annual income	1.349 (0.705)	2.244 (1.408)	1.655 (0.900)	2.335 (1.394)
Annual expenses	1.175 (0.624)	1.805 (1.222)	1.379 (0.797)	1.901 (1.133)
Total assets	1.076 (1.473)	3.954 (10.05)	1.454 (3.932)	4.056 (5.415)
Observations	2,900	2,357	1,218	973

Notes: This table reports means for the 1995 and 2002 CHIP data, by household type. School-age households report having at least one child aged 6-18 and none older. College-age households have at least one child aged 18-23. Standard deviations are in parentheses. Income, expenses, and assets are in 10,000s of yuan at current (unadjusted) prices.

but we also make use of households with college-age children. We define school-age as between 6 and 18 because it covers the usual elementary through high school years. We define college-age as between 18 and 23, and consider all households containing any college-age child as a college household. Households with children between 6 and 18 (but none between 18 and 23) count as school-age children households. We do not use households in which all children are older than 23 or younger than 6. To hone in on prime-age workers, we drop households whose head's age is less than 25 or above 60.

Our key outcome of interest is the saving rate, as calculated for each household using Eq. (1). We drop the few households that changed Hukou after the expansion but before their children took the college entrance exam. We also drop households from the 2002 sample in which a child took the exam before 1999 (when the expansion began) because we will use the households with college-age children to estimate the change in the likelihood of college enrollment (due to the policy). The final data set consists of 2900 school-age children households and 1218 college-age households in 1995 and 2357 school-age children households and 973 college-age households in 2002.

Table 1 reports summary statistics by household type (college or school-age children) for 1995 and 2002, using current (unadjusted) prices. The average saving rate for school-age children households rises considerably between 1995 and 2002, more than a 30 percent increase from 10.7 to 14.2. As the large standard deviations attest, there is ample variation in saving across households to exploit. In the remainder of the paper, we link the saving behavior to the expansion of higher education. Note, that over the same time period, the average saving rate among households with college-age children actually decreased. We do not directly study this decrease, but it might be related to the larger share of college-age children attending college in 2002 (18.3 percent) versus 1995 (10.1 percent).

We use this observed increase in college attendance to estimate the changes in the expectation of attending college for younger children.

The remainder of Table 1 reports the means and standard deviations for the control variables that we include in our regressions.⁹ Many of the controls also had large changes over time. For example, China's rapid growth pushed up incomes, assets, and expenses. Other changes can be traced to specific policies. The decreasing number of children was likely due to fertility policies set in the 1970s. While changes in home ownership and employment at State-Owned Enterprises (SOEs) can be traced to privatization policies enacted in the late 1990s (Berkowitz, Ma, & Nishioka, 2017; Chen et al., 2020; Chen & Wen, 2017). Importantly, we not only control for variables related to these other policy initiatives, but we also allow their impact on saving rates to vary over time. We also run several additional robustness checks to show that these other policies do not drive our results.

As mentioned, we use the sample of households with college-age children to estimate the change in college probability. However, since quotas for college admittance were set at the province level, we also use province specific information on enrollment rates. For each year and province, we approximate the enrollment rate with the ratio of new college students to the number of senior high school graduates.¹⁰ Even though it is possible that students go out of their home provinces for college, the vast majority of them remain in their home provinces. To support the idea that college enrollments are correlated with college probabilities, we supplement our analysis using the Census 2000 data. Appendix Table C1 shows the percentage of college students in a province that are from this province. We use number of college students from the province divided by total number of students attending college in that province. In most cases, more than 70% of college students in that province were originally from that province. Municipalities like Beijing have smaller shares due to a larger higher education system therefore they admit more students from other provinces. We address this issue in robustness checks by dropping households reside in Beijing. Using the Census 2000 data we find supportive evidence that the provincial level college enrollment rates are good predictors for college probabilities. We select school-age people (18–22) and merge college enrollment rates in the college-going year with each individual. Then we perform probit analyses with different specifications to see whether they can meaningfully determine the likelihood of going to college. The results are in Appendix Table C.2.¹¹ The estimates for provincial enrollment rates across columns are statistically significant, which provide a consistent story that enrollment rates at the province level can predict college status.

Fig. 5 shows the estimated college enrollment rates from the 12 provinces represented in our dataset from 1990 to 2002. In general, there is an upward trend, although many provinces dip just prior to the expansion due to the cohort effect where more high school students graduated starting in 1996. After 1998, enrollment rates increase for all provinces. But neither the enrollment rate levels nor the changes in enrollment rates are uniform across provinces. These differences in enrollment changes were driven by policy, and they therefore help us to identify the impact on savings. The next section (Eq. (11), in particular) provides further details.

5. Empirical approach

The theoretical model in Section 3 emitted a straightforward relationship between the policy-induced change in a household's expectation of college attendance and its saving rate. However, the available data (comprised of repeated cross-sections) does not track the same households over time, and it does not contain a measure of the expected probability of attending college. Ideally, we would estimate a difference-in-differences model where we would have an indicator for families who were affected by the policy. Some families were still not able to send their child to college (e.g. if their child still did not meet the academic criterion). Some families met the criterion for entrance prior to the expansion. Thus, our ideal specification would be a differences-in-differences specification:

$$SR_{ijt} = \alpha Post_t + \gamma Treat_{ij} + \pi Treat_{ij} Post_t + X'_{ijt} \delta + \lambda_j + u_{ijt}. \quad (7)$$

The dependent variable is the saving rate (SR_{ijt}) for household i from province j in year t . $Post_t$ is a dummy variable that equals one if the household is observed in year 2002 after the expansion, i.e., $Post_{2002} = 1$. $Treat_{ij}$ is also a binary variable that equals one if the household was affected by the policy. The vector X_{ij} not only contains the control variables listed in Table 1 but also a quadratic function of household income due to the nonlinear relationship between income and saving rate (Huggett & Ventura, 2000). Vector λ_j represents province fixed effects, and u_{ijt} captures measurement error. The coefficient γ measures any difference between the treated and control group prior to the treatment. The coefficient α measures a shift in savings before and after treatment not associated with treatment. The coefficient π is the treatment effect. Alternatively, estimates could be obtained by two regressions, one for each time period with the treatment indicator in each. The treatment effect is then the difference between the two coefficients on treatment (Post – pre). This latter specifications is algebraically equivalent to interacting the $Post$ variable with all X 's, and thus allows for differences in response to other variables.

The problem, however, is that we cannot directly identify, in families with school aged children, who would find their expectation of college attendance changed. Thus we do not have a treatment indicator readily measured. To address this we estimate models of college attendance using the sample of families with college-age children. The difference in the predicted probabilities serves as a measure of the treatment variable, we label as Δp_{ij} . We use the provincial enrollment rates for the families to identify the expansion

⁹ Appendix B provides definitions for each variable.

¹⁰ As far as we know, publicly available data on provincial college admission rates does not exist. Fan et al. (2018) uses university-specific cutoff scores collected from newspapers and websites to examine college expansion in China, and Hu & Bollinger (2017) use the number of people taking the college entrance exam and subsequent enrollment at the national level.

¹¹ Linear regressions have similar results.

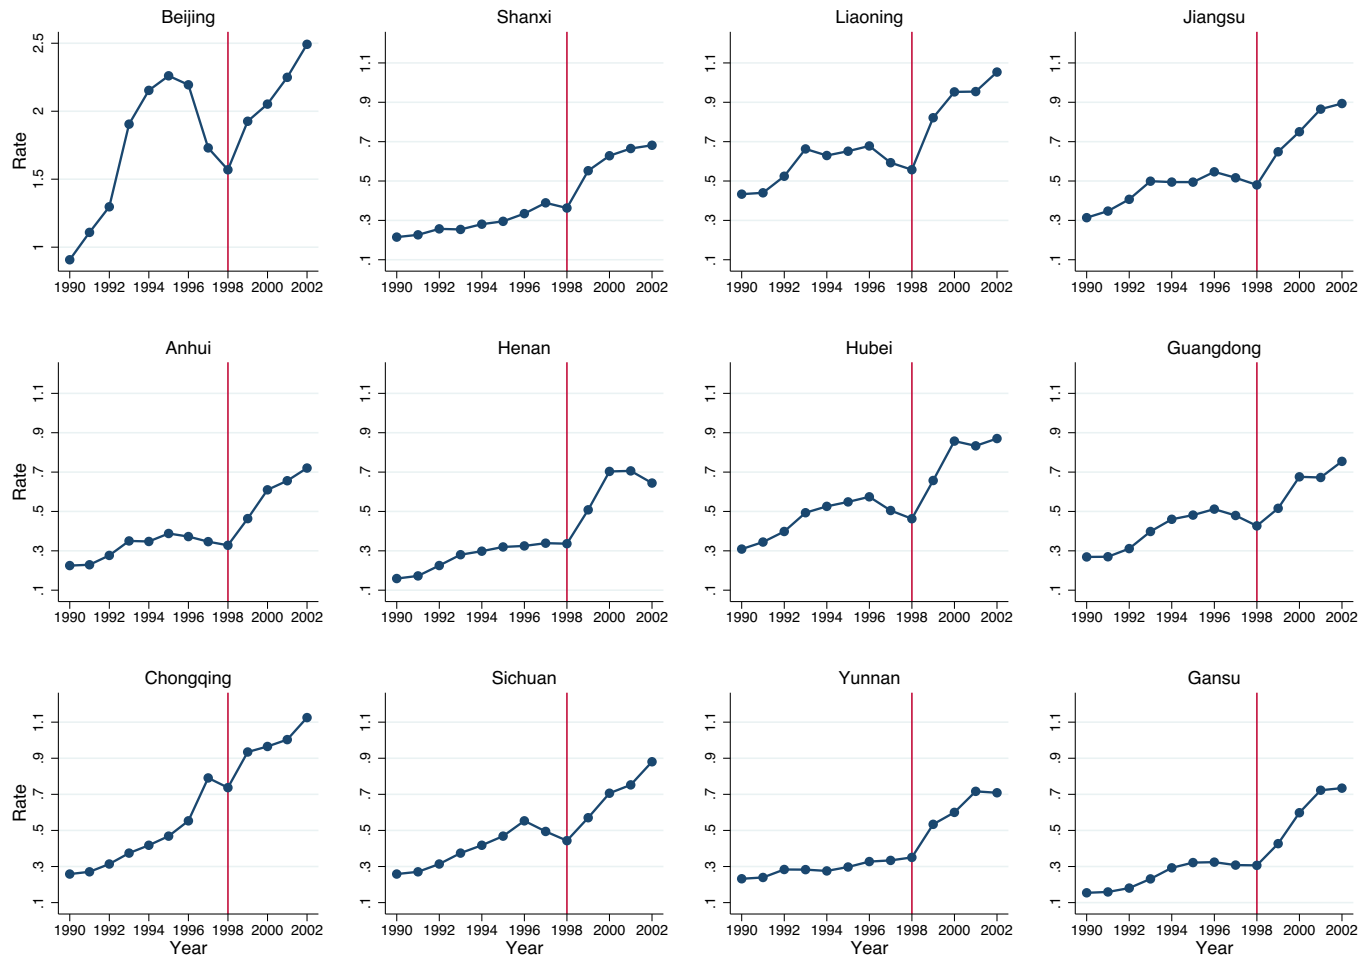


Fig. 5. Enrollment rate by province. *Notes:* We obtained the underlying data from the national bureau of statistics. This figure plots the annual rate of planned new enrollment in regular higher institutions divided by the number of senior high school graduates, from 1990 to 2002 by province. Prior to 1997, Chongqing was part of Sichuan, so we use the Sichuan rate.

impact. We then estimate the following Eq. (8) twice; once using the 1995 sample of households with school-age children and once using the 2002 sample of households with school-age children.

$$SR_{ij} = \pi \Delta p_{ij} + X'_{ij} \delta + \lambda_j + u_{ij}. \tag{8}$$

The coefficient (π) on the change in college probabilities (Δp) is our key parameter of interest. We obtain two estimates, one for 1995 (prior to the policy change) and one for 2002 (after the expansion). In each separate regression, π measures how much of a difference having a higher Δp makes for saving. The difference between the two estimates, $\pi^{2002} - \pi^{1995}$, then, is similar to a differences-in-differences estimate, with the difference attributable to the education expansion.

To see how this compares to a difference-in-differences estimator, we consider Eq. (7) in each of the two time periods. In the post period, Eq. (7) resolves to

$$SR_{ij1} = \alpha + (\gamma + \pi) \Delta p_{ij} + X'_{ij1} \delta + \lambda_j + u_{ij}. \tag{9}$$

Thus the coefficient on our proxy for treatment, Δp_{ij} , in the post period is $\gamma + \pi$. For the pre-period, Eq. (7) resolves to

$$SR_{ij0} = \gamma \Delta p_{ij} + X'_{ij0} \delta + \lambda_j + u_{ij}. \tag{10}$$

The coefficient corresponding to π is the difference between the coefficient on Δp_{ij} in each period. As noted above, we do not have the ideal treatment variable. We use the difference in probabilities as a measure of the likelihood that a particular household would be affected by the change in test score cutoff.

Note that the coefficients (δ and λ 's) are estimated separately for 1995 and 2002. The impact from the controls can vary over time.¹² Recall, for example, that the housing market changed during the late 1990s. Our regressions allow the relationship between home ownership and savings to change, accordingly. We return to housing (and other) reforms in Section 7.1. Alternatively we could have simply regressed the saving rates on the probability of attendance in each year, and differenced the coefficients. Unsurprisingly, the results are similar, however this is not exactly a difference-in-differences estimate, as it forces the coefficient on treatment to be zero prior to the treatment period. We prefer our approach, and focus on that methodology here. The alternative approach is available from the authors, and indicates a larger effect.

To construct Δp_{ij} we estimate Eq. (11) two times, once using households with college-age children in the 1995 sample and separately using the 2002 college-age households.

$$p_{ij} = X'_{ij} \beta + \gamma f(ER_i) + \zeta_j + \nu_{ij}. \tag{11}$$

The dependent variable (p_{ij}) equals one if the household's college-age child is enrolled in college, and zero otherwise. The share of college-age children attending college jumps from about 10 percent in 1995 to over 18 percent in 2002; see Table 1. The control variables are the same as above, and ζ represents a full set of province fixed effects. $f(ER_i)$ is a polynomial function of enrollment rates. Note that ER_i is the household specific enrollment rate corresponding to the year that the household's child took the college entrance exam in province j . The enrollment rate is not only different across provinces, but also different for households if their children's college-going years are different.

With estimates of β , γ , and ζ from both 1995 and 2002 in hand, the second step is to calculate the predicted probability of college enrollment, p , for each household in the subset of families with young (age 6 to 18) children. To do so, we plug in the observable data (X_{ij} and the province) into the estimated version of Eq. (11) to get an estimate of p for each household with young children.¹³ We do this twice for each household (and for both the 1995 and 2002 samples) – once using the vector of coefficient estimates based on 1995 college-age households and once using the estimates from 2002. We then use the difference in the two estimated probabilities as our measure of the household's change in the expected probability of sending their children to school. Specifically, ER_j^{2002} is assigned to households with school-age children in province j in 2002 and ER_j^{1995} is assigned to households with school-age children in province j in 1995.

$$\Delta p_{ij} = X'_{ij} \cdot (\hat{\beta}^{2002} - \hat{\beta}^{1995}) + (\hat{\gamma}^{2002} ER_j^{2002} - \hat{\gamma}^{1995} ER_j^{1995}) + (\hat{\zeta}_j^{2002} - \hat{\zeta}_j^{1995}). \tag{12}$$

We assume that the variation in Δp across households was exogenously driven by the higher education expansion. If this were strictly true, then we would only need to calculate Δp for the 2002 households and estimate Eq. (8) one time. However, the treatment intensity could have been correlated with unobservable household characteristics. Thus, we estimate Δp and π for the 1995 households in order to difference out any spurious (or pre-existing) correlation between household saving and Δp . In a sense, the Δp estimates for the 1995 households act as a placebo to check the efficacy of the treatment on the 2002 households.

Our other identifying assumption is that households experiencing different treatments (low changes in college probability versus high) would have had similar trends in their saving rates if the college expansion had not occurred. In Section 6.2, we use aggregate

¹² Thus, the use of micro data allows us to net out changes in the relationship between other observables and saving behavior by households with young children in a way that would not be possible with only aggregated provincial data. Relatedly, as we discuss below, we can estimate a household specific change in enrollment probability.

¹³ These children had not yet taken the college entrance exam, so we use the survey year enrollment rate for ER .

data to provide evidence that the saving rate trends were in fact similar across households in the years leading up to the college expansion.

6. Higher education expansion's effect on saving

This section reports our main findings, which are obtained by estimating Eq. (8) via ordinary least squares and Eq. (11) as a probit model. According to our estimates, a 10 percentage point increase in the expected probability of attending college leads to a 5.9 point increase in the average household's saving rate. The estimates are statistically significant and robust to a host of specification alterations. We begin by discussing the probit model.

6.1. Who goes to college?

The 1999 higher education expansion increased college opportunities for nearly all families. To calculate household-specific changes in the expectation of college attendance, we first estimate Eq. (11) using the subset of families with children in their college years (age 18 to 23), separately for 1995 and 2002. The dependent variable is the dummy variable indicating whether a household has a college child, and the controls include all those listed in Table 1, as described above. The Appendix reports the probit estimate details.

We then use the resulting coefficient estimates from Eq. (11) to calculate an expected college probability for each household with young children (ages 6 to 18). Fig. 6 shows the distribution of predicted college probabilities for the 2002 households using both sets of estimated coefficients. The distribution shifts to the right when using the 2002 coefficients. After the 1999 college expansion, the predicted probability of attending college increases markedly. The increase is mostly driven by the increase in provincial enrollment rates, but other factors matter, too.

Finally, we calculate the change in college probability using Eq. (12). Again, we do this for both the 1995 and 2002 households with young children. To get an initial sense of the relationship between the change in college probability and saving rates, we also regress the household specific saving rates on the controls listed in Table 1 (for the 2002 households). Fig. 7 plots a fitted curve (smoothed by a local polynomial regression (Fan & Gijbels, 1996) of the residuals against the change in college probability. The curve lines up closely with the predictions of our theoretical model. Small increases in college probability appear to be less correlated with saving. However, once the change becomes large enough, saving rates are higher, too. We next estimate this relationship using a quasi-difference-in-differences approach, in order to net out unobservable factors for households experiencing different changes in college probability.

6.2. The main empirical findings

Table 2 reports the main regression results based on Eq. (8). The estimate of π (the marginal effect of the change in college probability) for the 2002 sample equals 0.486. The effect on saving is quantitatively large and statistically different from 0 at better than the 5 percent significance level. Households for which the probability of attending college went up saved more. The estimate for the 1995 sample is much smaller (−0.108). Recall, the 1995 sample was not (yet) actually subject to the college expansion, while the 2002 sample was. Thus, we interpret the difference in the two coefficient estimates as the effect due to the policy change. This 'difference-in-differences' (DD) estimate equals 0.594, and it is statistically different from zero. Taken literally, the DD estimate implies that, on average, a 10 percentage point (policy driven) increase in the expected probability of college increased the typical households




Fig. 6. Estimated college probability for 2002 households. Notes: This figure shows the distribution of estimated college probabilities across the households surveyed in 2002, based on the coefficient estimates from both the 1995 and 2002 probit regressions of Eq. (11).




Fig. 7. Saving rate and change of probability for 2002 households. *Notes:* This figure shows the relationship between saving rates and the change in college probability for 2002 households via local polynomial regressions.

saving rate by about 5.9 percentage points. This effect is very large, but given the dramatic increase in saving observed in China after 2000, it is not implausible. Also note, the impact on aggregate saving is not straightforward to calculate because households with older children or no children likely reacted to the policy changes, possibly through general equilibrium effects on future employment or wages. The alternative specification is pooling the two years together and interact each variable with the year dummy including provincial variables such as population, urban employment, and GDP growth rate. We provide the result in Column (3) and it shows similar estimate compared to the more preferred flexible specification. [Chamon and Prasad \(2010\)](#) compare households with different age children and find that the effect of having a school age child increases saving rate by 2 percentage points before 1999 and 4 percentage points after 1999. The difference is approximately 2 percentage points. However, they do not utilize cross-province differences and examine the saving rate in an aggregate level.

[Table 3](#) reports two sets of robustness checks.¹⁴ All the regressions include the full set of controls, but we do not report the coefficient estimates to save space. In the first set (columns 1–4), we drop households residing in Beijing (columns 1 and 2) and Chongqing (columns 3 and 4). We omit Beijing because it has the largest higher education system in China. The overall enrollment rates exceed one which means it absorbs more college students than its own high school graduates. It can also be seen from the Appendix Table C1 that local students have much smaller percentage compared to other provinces. We omit Chongqing because by the year 1995 it was still part of Sichuan Province. It became a municipality in 1998, in order to ensure that our estimate is not driven by relative larger size of colleges after Chongqing became a municipality, we omit it to test the robustness. Dropping Beijing and Chongqing is somewhat arbitrary, but our estimates for the changes in college probability are largest for households in these two provinces, on average (see [Fig. 5](#) and Appendix Table C.4). The resulting DD estimates are smaller than our main DD estimates; however, the estimated effect is still very large and statistically significant. We retain the households from Beijing and Chongqing in the remainder of our analysis.

The second set of robustness checks adds households with zero dependent children in the sample. If it is the case that households with dependent children were expecting increases in college probability while in contrast households without children were not affected by such expansion, we should expect an even larger effect.¹⁵ Again, we find that the increase in college probability leads households to save more.¹⁶ We assign $\Delta p = 0$ to all households without dependent children, which has 686 households in 1995 and 633 households in 2002. The inclusion of this group results in an estimated impact of 8.34 percentage points, compared to our baseline result of 5.94. Overall, the results across the robustness checks reported in [Table 3](#) reinforce our main results. The number of observations is less than in the main regressions, but the effects remain statistically significant and quantitatively large. The empirical results in [Tables 2](#) and [3](#) are consistent with the intuition coming out of our structural model that low-saving households ($S < \tau$)

¹⁴ As an additional robustness check, we also have conducted a standard placebo check by randomly assigning households to different provincial enrollment rates. We find no significant impact on saving, providing additional support for our main results.

¹⁵ Age-savings profiles show supportive evidence for households without children do not respond to the college expansion. The figures are in Appendix E.

¹⁶ We also conducted the analysis using zero dependent child households as control group and households with child as treatment group. We found small positive impact. The results are available upon request.

Table 2
College probability's effect on household saving rates.

	(1)	(2)	(3)
	1995	2002	Pooled
Δp	-0.108 (0.077)	0.486** (0.241)	-0.068 (0.081)
<i>Post</i>			0.069 (0.481)
<i>Post</i> \times Δp			0.489** (0.240)
No. kids	0.008 (0.010)	0.014 (0.016)	0.006 (0.009)
No. elderly people	-0.004 (0.006)	0.006 (0.006)	-0.004 (0.005)
Private house	0.007 (0.006)	-0.036** (0.016)	0.006 (0.006)
Housing accumulation fund	0.002 (0.003)	-0.001 (0.012)	0.001 (0.003)
Head gender	0.005 (0.007)	0.028*** (0.008)	0.005 (0.006)
Currently working	0.008 (0.014)	0.022 (0.019)	0.008 (0.012)
Head tenure	0.000 (0.000)	0.000 (0.001)	0.000 (0.000)
Head age	-0.001 (0.000)	-0.001 (0.001)	-0.000 (0.000)
Head with college degree	-0.020 (0.018)	-0.023 (0.016)	-0.018 (0.016)
Head years of schooling	-0.000 (0.000)	0.001 (0.001)	0.000 (0.000)
Spouse years of schooling	0.000 (0.001)	0.005*** (0.002)	0.000 (0.001)
SOE job	0.008 (0.005)	0.025* (0.013)	0.007 (0.005)
Public health	-0.010 (0.008)	0.019* (0.010)	-0.007 (0.007)
Annual income	0.602*** (0.093)	0.359*** (0.077)	0.602*** (0.080)
Annual income square	-0.028* (0.015)	-0.011 (0.009)	-0.028** (0.012)
Annual expense	-0.551*** (0.084)	-0.333*** (0.038)	-0.549*** (0.075)
Total asset	-0.001 (0.002)	-0.003** (0.001)	-0.001 (0.002)
Population, millions			-0.009 (1.614)
Urban employment, millions			0.119 (1.048)
GDP growth rate			-0.023 (2.745)
Observations	2,900	2,357	5,257
R^2	0.787	0.713	0.739
DD		0.594 (0.253)	

Notes: This table reports regression estimates of Eq. (8) for 1995 and 2002. Income, expenses, and assets are in 10,000s of yuan and current prices. Row DD reports the difference in the coefficient estimates (π) on Δp . The regressions include province fixed effects, and the parentheses report bootstrapped robust standard errors clustered at province level with 500 repetitions.

experiencing a large enough increase in the likelihood of college (p) will start to save more (i.e. Proposition 1).

We end this section by revisiting our identification strategy. A critical assumption for the difference-in-differences framework is that the differentially treated groups were not trending (before the treatment) in ways correlated with the treatment.¹⁷ In our case, the concern is that the types of households experiencing large increases in their college probability were also the types of households

¹⁷ Another critical assumption is that households did not select into experiencing a higher increase in college probability, for example, by moving across provinces. This is unlikely to have occurred because of the speed of the policy implementation and the Hukou restrictions on migration. Just to be sure, we have checked that our results do not change when we remove the few households in the sample that changed their Hukou after the policy.

Table 3
Robustness checks.

	No Beijing		No Chongqing		Add zero child	
	(1)	(2)	(3)	(4)	(5)	(6)
	1995	2002	1995	2002	1995	2002
Δp	-0.130 (0.137)	0.585*** (0.215)	-0.106 (0.077)	0.553** (0.262)	0.065 (0.128)	0.899*** (0.311)
Observations	2,690	2,225	2,900	2,268	3,586	2,987
R^2	0.780	0.709	0.787	0.710	0.773	0.380
DD		0.715 (0.254)		0.659 (0.273)		0.834 (0.318)

Notes: Each regression is based on Eq. (8). The dependent variable is saving rate. Each regression includes the full set of controls and province fixed effects. The parentheses report bootstrapped robust standard errors clustered at province level with 500 repetitions.

already increasing their savings, anyway. This concern is valid, and we can even imagine that the policy makers might have wanted to target the treatment towards households with increasing rates of savings because these households would be most able to pay for college. With the cross-sectional data at hand, we cannot directly check the micro-data for pre-trends; however, we can take a more aggregated look at the data. Fig. 8 plots changes in the average household saving rate before the education expansion (from 1995 to 1998) against enrollment rates during the enrollment increase (1998 to 2002), for each province. There exists little correlation, and the fitted line actually reveals a slightly negative relationship. In other words, the policy did not selectively increase enrollment rates more for the provinces in which saving rates were already growing the most before the policy was enacted. Recall that the household-level estimates for the change in college probability (Eq. (11)) were a function of the provincial enrollment rates. So, from this perspective, the treatment on individual households does not seem to depend on their pre-policy saving behavior. Thus, we conclude that our analysis is not picking up pre-trends, but, instead, the saving behavior is driven by the response to changing college opportunities.

7. Other explanations for China's high saving rates

This section examines the other main explanations for China's high household saving rates put forth in the literature (see Section 2). First, we examine policy changes enacted around the same time. Then, we examine demographic changes. As noted above, our regressions already include a number of controls aimed at accounting for these related factors.¹⁸ Moreover, we think these other factors were unlikely to be linked to the pattern of increased college opportunities. However, since the policy reforms and demographic changes were so big and also impacted household saving, we now go beyond simple control variables and allow for interactions with our variable of interest. Throughout the many specifications, the impact of higher education expansion on saving rates remains large and statistically significant.¹⁹

7.1. Other reforms

We begin by interacting the change in college probability (allowing for heterogeneous effects) with the variables related to three other large reforms: SOE reform, housing reform, and healthcare reform. Specifically, we estimate

$$SR_{ij} = W'_{ij}\delta + \pi \cdot \Delta p_{ij} + \varphi \cdot Z_{ij} + \psi \cdot \Delta p_{ij} \cdot Z_{ij} + \lambda_j + u_{ij}, \quad (13)$$

where Z_{ij} represents a dummy variable capturing either SOE job status, public health coverage, or private house ownership (each is considered separately). Vector W is the vector X , including all covariates except Z .²⁰ Then, $\partial SR_{ij}/\partial \Delta p = \pi + \psi \cdot Z_{ij}$ is the marginal effect for each household. We report the average marginal effects, based off of the marginal effects across different Z_{ij} 's. We are interested in whether controlling for the heterogeneous impact across Z (i.e. the exposure to other reforms) alters our estimate (π) for the education expansion effect, and we find that it does not.

7.1.1. Employment at a state-owned enterprise (SOE)

Public sector jobs used to be part of the so-called "iron rice bowl" of social support. The SOE reform in the 1990s, however, led to massive layoffs and made even incumbent SOE workers less secure in their employment. As discussed above, several papers have shown that the new risk led to an increase in precautionary saving. If the households that experienced large increases in the likelihood of college were also facing greater employment risk, then our estimates could be conflating the two channels. Columns (1) and (2) of Table 4 provide evidence against this possibility. The coefficient estimate for Δp (0.424) in 2002 is close to the main estimate, even

¹⁸ Our main regressions also allow the coefficient estimates on these controls to change over time.

¹⁹ The Appendix also reports triple difference-in-difference estimates, which leverage these other factors affecting saving rates to further difference out potential trends coming from a third, omitted, dimension (Imbens & Wooldridge, 2009).

²⁰ To save space, the tables below do not report the coefficient estimates for the controls.

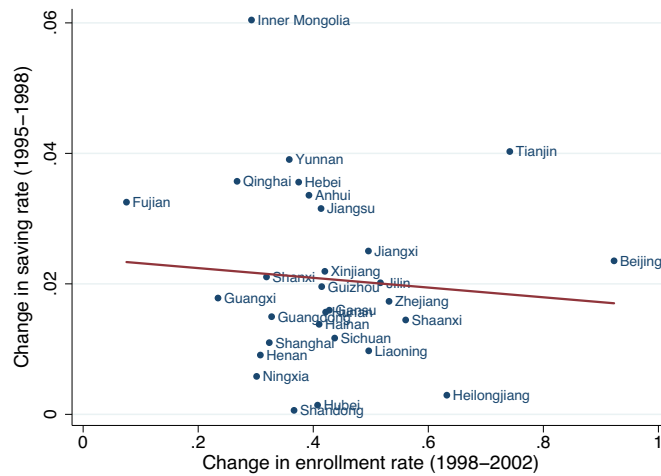


Fig. 8. Test for pre-trends. *Notes:* The vertical axis measures the saving rate change between 1995 and 1998. The horizontal axis measures the subsequent enrollment rate changes between 1998 and 2002. The underlying data comes from the national bureau of statistics yearly statistical book.

Table 4
Other reforms.

	SOE job		Public health		Private house	
	(1)	(2)	(3)	(4)	(5)	(6)
	1995	2002	1995	2002	1995	2002
Δp	-0.132 (0.164)	0.424** (0.216)	-0.146 (0.158)	0.521** (0.248)	-0.092 (0.072)	0.477** (0.235)
$\Delta p \times$ SOE job	0.023 (0.148)	0.187 (0.158)				
$\Delta p \times$ Public health			0.036 (0.123)	-0.035 (0.132)		
$\Delta p \times$ Private house					-0.124 (0.147)	0.010 (0.160)
SOE job	0.008 (0.007)	0.029** (0.013)	0.008 (0.005)	0.025* (0.013)	0.008 (0.005)	0.025* (0.013)
Public health	-0.010 (0.008)	-0.019* (0.010)	-0.009 (0.010)	0.018 (0.011)	-0.010 (0.008)	0.019* (0.010)
Private house	0.007 (0.007)	-0.036** (0.016)	0.007 (0.007)	-0.037** (0.016)	0.002 (0.007)	-0.036* (0.018)
Observation	2900	2357	2900	2357	2900	2357
R-square	0.787	0.713	0.787	0.713	0.787	0.713
Marginal effect	-0.113 (0.080)	0.484 (0.240)	-0.120 (0.092)	0.498 (0.235)	-0.143 (0.104)	0.485 (0.239)
DD		0.597 (0.253)		0.618 (0.252)		0.628 (0.260)

Notes: Each regression includes province fixed effects and the control variables. The parentheses report bootstrapped robust standard errors clustered at province level with 500 repetitions.

after controlling for the interaction with SOE status. Note, the interaction term is positive in both years, but small and statistically insignificant. The bottom line is that the DD estimate (5.97 percentage points) remains large and close to our baseline result.

7.1.2. Public health coverage

Healthcare is another motive for precautionary saving (Chamon & Prasad, 2010), as well as life-cycle saving. Access to public health coverage has been in flux in China due to policy changes, migration and Hukou regulations, and population aging. Columns (3) and (4) of Table 4 report estimates controlling for heterogeneous effects based on having public health coverage. The resulting DD estimate of 0.618 is actually slightly higher than our baseline estimate. Interestingly, the coefficient estimate for the interaction term flips signs and becomes statistically insignificant after the college expansion.

Table 5
Demographics.

	Number of kids		Gender of kids		Life-cycle	
	(1)	(2)	(3)	(4)	(5)	(6)
	1995	2002	1995	2002	1995	2002
Δp	-0.429 (0.331)	0.602* (0.347)	-0.174 (0.130)	0.478* (0.247)	-0.040 (0.577)	0.929* (0.501)
Single	0.032 (0.060)	0.003 (0.058)				
$\Delta p \times$ Single	0.340 (0.309)	-0.122 (0.247)				
Male			0.004 (0.003)	0.003 (0.008)		
$\Delta p \times$ Male			0.110 (0.125)	0.012 (0.124)		
Head age	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.000)	-0.001 (0.001)	-0.001 (0.001)	-0.002 (0.001)
$\Delta p \times$ Head age					0.004 (0.011)	-0.029 (0.009)
Observation	2,900	2,357	2,900	2,357	2,900	2,357
R^2	0.787	0.713	0.787	0.713	0.787	0.713
Marginal effect	-0.126 (0.086)	0.490 (0.249)	-0.121 (0.086)	0.484 (0.241)	-0.096 (0.140)	0.515 (0.249)
DD		0.616 (0.2638)		0.605 (0.256)		0.611 (0.286)

Notes: Each regression includes province fixed effects and the control variables. The parentheses report bootstrapped robust standard errors clustered at province level with 500 repetitions.

7.1.3. Home ownership

China also reformed its housing market in the 1990s. By 1998, most families were allowed to own their homes, and many households saved in the form of housing or in order to buy housing (required down payment rates were high). Again, as with the other reforms, we think that a connection between housing reform and exposure to the higher education expansion is unlikely at the household level. Columns (5) and (6) in Table 4 indicate that controlling for the interaction between the change in college probability and home ownership has little impact on our main estimate. The DD estimate (0.628) remains very large.

7.2. Household demographics

We next consider regressions that control for interactions with variables capturing household demographic characteristics. We still apply the approach embodied in Eq. (13), letting Z represent the various demographic factors of interest.

7.2.1. Number of children

China enacted the One Child Policy in 1978, although enforcement of the policy varied over time and by location. Several papers have shown that Chinese households with fewer children (i.e. one) tend to save more and have drawn a connection between China's fertility policies and high saving rates. One reason why single child households might save more is to invest in their child's college education (a quantity/quality tradeoff similar to Becker & Lewis (1973)). In order to examine the potential heterogeneous effect across households with different numbers of children, we define a new variable *Single* that is equal to one if there is only one dependent child and zero if there are two or more children. We re-estimate Eq. (13), interacting the new variable ($Z = \text{Single}$) with the change in the probability of college.

Table 5 columns (1) and (2) present the results. The resulting DD estimate for the average marginal effect equals 0.616, close to our baseline result. Note, the coefficient estimate on *Single*, while not statistically significant before the education expansion, is close to what Lugauer et al. (2019) estimated, a 2.4 percentage point decrease in saving rates due to each additional child using a different data set and different estimation methodology.

7.2.2. Sex composition

Some Chinese households may prefer having a son over a daughter because sons traditionally support their parents in old age. Relatedly, the One Child Policy may have led to the gender imbalance now prevalent in China. Households with sons invest differently for their child's education and also for marriage and housing purposes. Wei and Zhang (2011) show that these factors have had a large impact on household saving rates in China.²¹ Therefore, we next interact college probability with a variable ($Z = \text{Male}$) that equals one

²¹ Cai et al. (2019) and Zhou (2014) also discuss related issues.

if a household has only male children and is zero otherwise. We count mixed gender households (of which there are few) as non-male households. Columns (3) and (4) of Table 5 suggest that the male child effect is quite small. Moreover, our estimate of the college expansion effect (DD=0.605) remains large.

7.2.3. Life-cycle saving and age effects

Simple life-cycle theory predicts that households of different ages can have different saving patterns. The younger households in our sample (household head with an age closer to 25) might be net borrowers, while older households (age close to 60) rapidly accumulate assets in anticipation of retirement. The older the household head, the closer the child is to college age, on average. To address this, we next interact the household head's age ($Z = \text{Age}$) with the change in college probability. Columns (5) and (6) of Table 5 present the results. Once again, the DD estimate (0.611) remains similar to our baseline result.

In summary, the 1999 higher education expansion had a large effect on household saving rates, and this finding holds across households with different exposures to other reforms and different demographic characteristics. While the economic reforms and demographic factors likely also impacted household saving (our regressions do not rule out these alternative stories), the increase in college opportunities had an additional effect above and beyond these other explanations.

8. Conclusion

In this paper, we exploit the policy-induced increase in college enrollment to estimate the impact on household saving rates in China. We find that the expansion in higher education resulted in higher household saving rates, especially for previously low-saving households experiencing a large increase in the probability of sending their children to college. Through this saving channel, the college expansion likely affected China's economic growth rate and international capital flows. Our findings are robust to a host of specification modifications, as well as to controlling for other concurrent policy changes and on-going demographic changes.

China's education expansion was unique in several ways. The policy change was large and swiftly implemented, and college enrollment levels were relatively low before the reforms. Also, China had strict fertility controls at the time, effectively shutting down the quantity channel in the classic fertility theory of a quantity/quality trade-off. These characteristics helped inform our estimation strategy, but it remains a question as to whether our findings are applicable to other countries and situations. Similarly, our regressions necessarily capture the short run responses, from when the policy was new and unexpected, rather than the long-run response. We leave an exploration of these important issues to future research. Understanding the link between saving and education is important and informative for policy makers for many reasons. First, understanding saving behavior itself is, of course, important, and education expenses are a key part of household expenditures, over 10 percent on average in China. Second, many countries including the U.S. are debating whether to increase access to higher education, and it is natural to question how the additional schooling will be financed. Research is needed on how the expansion of education affects the economy and how households respond to the increase in college opportunities. Our story is straightforward. Households save to pay for college. Expanding higher education means that more families will expect their children to attend university; hence, these households save more.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.chieco.2021.101736>.

References

- Bairoliya, N., Miller, R., & Saxena, A. (2018). *The macroeconomic impact of fertility changes in an aging population*, Working paper.
- Banerjee, A., Duflo, E., Glennerster, R., & Kinnan, C. (2015). The miracle of microfinance? Evidence from a randomized evaluation. *American Economic Journal: Applied Economics*, 7(1), 22–53.
- Becker, G. S., & Lewis, H. G. (1973). On the interaction between the quantity and quality of children. *Journal of Political Economy*, 81(2, Part 2), S279–S288.
- Berkowitz, D., Ma, H., & Nishioka, S. (2017). Recasting the iron rice bowl: The reform of China's state-owned enterprises. *Review of Economics and Statistics*, 99(4), 736–747.
- Brown, M., Scholz, J. K., & Seshadri, A. (2012). A new test of borrowing constraints for education. *Review of Economic Studies*, 79(2), 511–538.
- Cai, X., Lu, Y., Pan, J., & Zhong, S. (2019). Gender gap under pressure: Evidence from China's national college entrance examination. *Review of Economics and Statistics*, 101(2), 249–263.
- Cameron, S. V., & Taber, C. (2004). Estimation of educational borrowing constraints using returns to schooling. *Journal of Political Economy*, 112(1), 132–182.
- Carniero, P., & Heckman, J. J. (2009). The evidence on credit constraints in post-secondary schooling. *The Economic Journal*, 112(482), 705–734.
- Chamon, M. D., & Prasad, E. S. (2010). Why are saving rates of urban households in China rising? *American Economic Journal: Macroeconomics*, 2(1), 93–130.
- Chamon, M. D., Liu, K., & Prasad, E. S. (2013). Income uncertainty and household savings in China. *Journal of Development Economics*, 105, 164–177.
- Chao, C. C., Laffargue, J. P., & Yu, E. (2011). The Chinese saving puzzle and the life-cycle hypothesis: A reevaluation. *China Economic Review*, 22(1), 108–120.
- Chen, B., & Yang, R. (2012). Understanding China's rising saving rate: The role of higher education reform. In *The new Chinese economy* (pp. 137–153).
- Chen, B., Yang, X., & Zhong, N. (2020). Housing demand and household saving rates in China: Evidence from a housing reform. *Journal of Housing Economics*, 49, 101693.
- Chen, K., & Wen, Y. (2017). The great housing boom of China. *American Economic Journal: Macroeconomics*, 9(2), 73–114.
- Che, Y., & Zhang, L. (2018). Human capital, technology adoption and firm performance: Impacts of China's higher education expansion in the late 1990s. *The Economic Journal*, 128(614), 2282–2320.
- Chi, W., & Qian, X. (2016). Human capital investment in children: An empirical study of household child education expenditure in China, 2007 and 2011. *China Economic Review*, 37, 52–65.
- Choi, H., Lugauer, S., & Mark, N. C. (2017). Precautionary saving of Chinese and U.S. households. *Journal of Money, Credit and Banking*, 49(4), 635–661.

- Choukhmane, T., Coeurdacier, N., & Jin, K. (2018). *The one-child policy and household savings*. Working paper.
- Curtis, C. C. (2016). Economic reforms and the evolution of China's total factor productivity. *Review of Economic Dynamics*, 21, 225–245.
- Curtis, C. C., Lugauer, S., & Mark, N. C. (2015). Demographic patterns and household saving in China. *American Economic Journal: Macroeconomics*, 7(2), 58–94.
- Curtis, C. C., Lugauer, S., & Mark, N. C. (2017). Demographics and aggregate household saving in Japan, China, and India. *Journal of Macroeconomics*, 51, 175–191.
- Fan, E., Meng, X., Wei, Z., & Zhao, G. (2018). Rates of return to four-year university education: An application of regression discontinuity design. *The Scandinavian Journal of Economics*, 120, 1011–1042.
- Fan, J., & Gijbels, I. (1996). *Local polynomial modelling and its applications: Monographs on statistics and applied probability* (Vol. 66). CRC Press.
- Feng, S., & Xia, X. (2018). Endogenous skill-biased technology adoption: Evidence from China's college enrollment expansion program. *Human Capital and Economic Opportunity Working Papers*, 2018–2099.
- Ge, S., Yang, D. T., & Zhang, J. (2018). Population policies demographic structural changes, and the Chinese household saving puzzle. *European Economic Review*, 101, 181–209.
- He, H., Huang, F., Liu, Z., & Zhu, D. (2018). Breaking the “iron rice bowl”: evidence of precautionary savings from the Chinese state-owned enterprises reform. *Journal of Monetary Economics*, 94, 94–113.
- Hu, C., & Bollinger, C. (2017). *China's higher education expansion and college premium: IV estimations*. Working Paper.
- Hu, C., & Bollinger, C. (2021). Effects of cohort size on college premium: Evidence from China's higher education expansion. *China Economic Review*, 70, 101700.
- Huggett, M., & Ventura, G. (2000). Understanding why high income households save more than low income households. *Journal of Monetary Economics*, 45(2), 361–397.
- Imbens, G. W., & Wooldridge, J. M. (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, 47(1), 5–86.
- Imrohorglu, A., & Zhao, K. (2018a). The Chinese saving rate: Long-term care risks family insurance, and demographics. *Journal of Monetary Economics*, 96, 33–52.
- Imrohorglu, A., & Zhao, K. (2018b). Intergenerational transfers and China's social security reform. *Journal of the Economics of Ageing*, 11, 62–70.
- Jia, N., Lugauer, S., Ni, J., & Yin, Z. (2021). *The effects of relaxing the one-child policy*. Working paper.
- Keane, M. P., & Wolpin, K. I. (2001). The effect of parental transfers and borrowing constraints on educational attainment. *International Economic Review*, 42(4), 1051–1103.
- Knight, J., Deng, Q., & Li, S. (2017). China's expansion of higher education: The labour market consequences of a supply shock. *China Economic Review*, 43, 127–141.
- Li, H., Ma, Y., Meng, L., Qiao, X., & Shi, X. (2017). Skill complementarities and returns to higher education: Evidence from college enrollment expansion in China. *China Economic Review*, 46, 10–26.
- Li, S., & Xing, C. (2010). *China's higher education expansion and its labor market consequences*. IZA Discussion Paper, 4974.
- Lugauer, S., & Mark, N. C. (2013). *Global Interdependence, Decoupling, and Recoupling* (p. 131). The MIT Press.
- Lugauer, S., Ni, J., & Yin, Z. (2019). Chinese household saving and dependent children: Theory and evidence. *China Economic Review*, 57, 101091.
- Malkova, O., & Braga, B. (2018). *Hope for the family: The effects of merit aid programs on parental outcomes*. Working paper.
- Modigliani, F., & Cao, S. L. (2004). The Chinese saving puzzle and the life-cycle hypothesis. *Journal of Economic Literature*, 42(1), 145–170.
- Rosenzweig, M. R., & Zhang, J. (2014). *Co-residence life-cycle savings and inter-generational support in urban China*. NBER Working Paper.
- Shen, H., & Li, W. (2003). *A review of the student loans scheme in China*. Bangkok: UNESCO Bangkok.
- Song, Z., & Yang, D. T. (2010). Life cycle earnings and the household saving puzzle in a fast-growing economy. *Manuscript*.
- Song, Z., Storesletten, K., Wang, Y., & Zilibotti, F. (2015). Sharing high growth across generations: Pensions and demographic transition in China. *American Economic Journal: Macroeconomics*, 7(2), 1–39.
- Sun, S. T., & Yannelis, C. (2016). Credit constraints and demand for higher education: Evidence from financial deregulation. *Review of Economics and Statistics*, 98, 12–24.
- Wan, Y. (2006). Expansion of Chinese higher education since 1998: Its causes and outcomes. *Asia Pacific Education Review*, 7(1), 19–32.
- Wang, X., Fleisher, B. M., Li, H., & Li, S. (2014). Access to college and heterogeneous returns to education in China. *Economics of Education Review*, 42, 78–92.
- Wei, S.-J., & Zhang, X. (2011). The competitive saving motive: Evidence from rising sex ratios and savings rates in China. *Journal of Political Economy*, 119(3), 511–564.
- Yao, R., Wang, F., Weagley, R. O., & Liao, L. (2011). Household saving motives: Comparing American and Chinese consumers. *Family and Consumer Sciences Research Journal*, 40(1), 28–44.
- Yoo, K., & Giles, J. (2007). Precautionary behavior and household consumption and savings decisions: An empirical analysis using household panel data from rural China. *Review of Economics and Statistics*, 89(3), 534–551.
- Zhou, W. (2014). Brothers, household financial markets and savings rate in China. *Journal of Development Economics*, 111, 34–47.