The Expansion of Higher Education and Household Saving in China

Xiaozhou Ding* Steven Lugauer† Christopher Bollinger‡

August 2019

Abstract

We investigate whether the expansion of Chinese higher education explains the increase in Chinese household saving. We motivate the work by a simple 2-period model of household saving for college. Using the China Household Income Project (CHIP) data, we estimate the change in the probability of college admittance following China's higher education expansion. We estimate how these changes in college probability affected household saving rates before and after the reform. We find that a 10- point increase in the probability of going to college raises the saving rate for a household by more than 7 percentage points.

Keywords: China; Higher Education Expansion; Saving Rate

JEL Classification Numbers: D12, D91, E21, I22, J13

*Corresponding author. Department of Economics, University of Kentucky, Lexington, KY 40506. Email: xiaozhou.ding@uky.edu.

†Department of Economics, University of Kentucky, Lexington, KY 40506. Email: steven.lugauer@uky.edu.

‡Department of Economics, University of Kentucky, Lexington, KY 40506. Email: crboll@email.uky.edu.

1 Introduction

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. The saving has helped maintain a high investment rate within China and allowed funds to flow abroad, particularly to the United States. Several theories have been put forward to explain the high saving rates, including demographic changes (Modigliani and Cao, 2004; Curtis et al., 2015; Imrohoroglu and Zhao, 2018b; Ge et al., 2018), income uncertainty (Chamon et al., 2013), private expenditures (Chamon and Prasad, 2010), economic reforms (He et al., 2018), and gender-related issues (Wei and Zhang, 2011; Zhou, 2014). None of the theories fully explains the changes. In this paper, we find evidence that the saving increase after 2000 is partly related to the expansion of college education.

Entrance into higher education in China is determined by an entrance exam and quotas set by the Ministry of Education. In 1999, the Ministry of Education began an extensive increase in the number of students allowed to enroll in college. The enrollment rate for high-school graduates quickly moved from approximately twenty percent to almost sixty percent, and is now eighty percent. China currently has almost forty million college students, about one-fifth of the college students in the world. This unprecedented education expansion likely impacts society in a number of ways. One of the most obvious effects, and the focus of this paper, is that Chinese households must finance their child's education. Since student loans were not common in China in the early years of higher education expansion (Shen and Li, 2003)[1], households pay for tuition out of accumulated savings or current income. Our hypothesis is that as enrollment rates increased, households with young children saved more in anticipation of future education related expenses.

Understanding the link between saving and education is important 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, researchers and policy makers are interested in explaining China's high saving rate in particular because, as mentioned, Chinese savings are a major component of international capital flows. China owns a large portion

[1] Only 32 percent of loan applicants succeeded in signing loan contracts and the loan coverage rate is less than 3 percent in 2001.

of European and US debt. The saving has also helped to fuel China's dramatic economic growth. Third, the relationship between saving and education may be informative for public policies of human capital accumulation. Many countries 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. We find that affected households sharply increase their saving.

To motivate our empirical approach and to clearly articulate the connection between education expansion and saving, we first present a simple two-period theoretical model of household saving. In the model, households receive income only in period 1 and must save for period 2. Households also can save to pay for college. If a household saves enough to pay tuition and expenses, and their child gets admitted, then the household reaps the benefits of having a college graduate. If a household does not save enough, their child cannot attend college. Thus, households compare the utility obtained from consuming more in period 1 against the expected utility from possibly sending their child to college and having greater consumption in period 2. As the expansion of higher education increases the probability of acceptance into college, the expected utility of saving for college increases. If the expected utility of saving for college becomes larger than that of not doing so, then the household increases their saving rate in order to have enough to pay for college. Thus, the model suggests that household saving rates should go up for households experiencing a large increase in the likelihood of sending their children to college.

Based on this intuition, we take 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. Our empirical strategy consists of two steps. The first step is to calculate the change in the expected probability of college for individual households. To do this, we estimate a probit model for the probability of children's college status using cross-sectional household level information from the 1995 and 2002 China Household Income Project (CHIP) surveys. The data allows us to see whether a family has children, their ages, and whether the children attend college. We estimate the probit model separately for each year using a host of control variables and only the subset of families with college-age children (households with a child aged 18 to 23). Then, we use the coefficient estimates on the controls to calculate the expected probability of college for the subset of families with children younger than college age (families with all their children aged 6 to 18). We do this twice for each household. Once with the coefficient estimates from the 1995 probit and once with the 2002 coefficient estimates. The difference between the two probabilities proxies for the change in the expected likelihood that a family's child will be able to attend college in the future. While college opportunities increased for almost all groups, the increase was highest for wealthy households with low levels of (parental) education from certain provinces. We assume that the policy (which induced the changes in college probability) was exogenous relative to household saving rates, and we provide province-level evidence in favor of this assumption.

In the second step of our empirical strategy, we estimate how the changes in college probability relate to household saving rates via a quasi difference-in-differences framework. Specifically, we regress household saving rates on a set of covariates and the estimated change of college probability (separately) for both households in the 1995 survey and 2002 survey. We include province fixed effects and allow all the parameter estimates (on the fixed effects and other controls) to differ between the 1995 and 2002 regressions to allow for variation in the observables' effects. We interpret the difference between the estimated coefficients on the change in college probability between the two years as the policy's impact on saving rates. We subtract the estimated coefficient estimate generated using the 1995 sample (before the policy was implemented) from the 2002 coefficient estimate in order to net out any time-invariant unobservables related to the intensity of treatment (i.e. the change in college probability). Given that the regressions include many controls, it is difficult to think what these unobservables might be. Nevertheless, our difference-in-differences approach helps to mitigate concerns over such omitted variables.

Consistent with the theoretical model, our main finding is that the change in the expected probability of college status has a quantitatively large and statistically significant effect on household saving rates. A 10 percentage point increase in college probability increases the average household's saving rate by 7.5 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 households that already had ample savings do not increase their saving rates; both these findings are also consistent with the theoretical model.

We report on a number of alternative specifications, and our main findings remain qualitatively unchanged. The robustness checks include omitting the provinces with the largest changes in enrollment rates and using wealth in place of savings. We also carefully consider other policy changes that occurred in China, including housing reform, public healthcare reform, and the reorganization of State-Owned Enterprises. Finally, we consider some of the large demographic changes that have been on-going in China, as previous literature has shown that demographic factors have impacted household saving in China. We discuss the burgeoning literature seeking to explain China's high household saving rate in the next section.

In addition to the recent research on China's high saving rates, our paper also is related to the literature on 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 et al. (2017), Li et al. (2017), and Bollinger and Hu (2017, 2018). These papers primarily focus on labor market implications, while we connect college expansion to saving behavior. Finally, our paper contributes to the growing literature that studies how decisions made within the family impact the macro-economy. See Doepke and Tertilt (2016), Doepke and Zilibotti (2019), and Greenwood et al. (2017) and the many recent papers cited within for more on this growing sub-field.

The rest of our paper is organized as follows. Section 2 documents the high 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

$$\text{Saving rate} = \frac{\text{Income} - \text{Expenditure}}{\text{Income}} = 1 - \frac{\text{Expenditure}}{\text{Income}}. \tag{1}$$

Figure 1 plots the urban household saving rate in China from 1981 to 2011. Saving slowly trended up before 2000; after that, though, the saving rate exploded. In recent years, the typical Chinese household has been saving 30 percent of their income, a truly astonishing number. This saving behavior is naturally of importance to Chinese households, but it also has important implications for the aggregate economy. The savings has contributed to the high investment rate (and therefore growth rate) in China. The excess saving has been flowing abroad, especially to the US and Europe. Therefore, 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.[2] As an example of this, He et al. (2018) show that the reform of state-owned enterprises in the 1990s increased unemployment and other risks for Chinese households, generating a dramatic increase in saving. Similarly, Choi et al. (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.[3] Several papers have documented specific examples of how demographics impact saving in China. Imrohoroglu and Zhao (2018a) argue that increases in

[2] 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).

[3] Also, see Bairoliya et al. (2018), Banerjee et al. (2015), Chao et al. (2011), Choukhmane et al. (2017), Curtis et al. (2017), and Lugauer et al. (2017).

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.[4] 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. 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.

Figure 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. Figure 3 plots this unprecedented increase in college enrollment alongside the national enrollment rate. The jump in the enrollment rate after 1999 (from below 40

[4]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 et al. (2016).

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 sizeable reduction in 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.[5] 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.

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.[6] Figure 4 plots the total annual tuition and fees per student 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 and 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.

[5] 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.

[6] 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 et al. (2012), Sun and Yannelis (2016) and Malkova and Braga (2018) for example.

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.

$$U = \ln C + \mathrm{E}[\ln C']. \tag{2}$$

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

$$C + S = Y, \tag{3}$$

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. 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

$$C' = S + I_c(\theta - \tau), \tag{4}$$

where θ is the benefit from sending a child to college and I_c equals one if the child attends college and zero otherwise.

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. First, 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.

Proposition 1 (Threshold p) *For a household choosing to save less than τ when $p = 0$, 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:*

$$p^t = \frac{\ln\left[\frac{Y^2}{4(Y-\tau)\tau}\right]}{\ln\frac{\theta}{\tau}}. \tag{5}$$

Proposition 1 has several implications.[7] Most relevantly, 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. This theoretical result helps to motivate our main empirical strategy. Households experiencing a large increase p are particularly likely to increase their saving to pay for tuition. We make use of this intuition by estimating the idiosyncratic increases in p across households and quantifying how changes in p impact saving. Proposition 1 also indicates 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. See the Appendix for more details.

Not all households react to the increase in p in the same way, though. Households experiencing a smaller increase in p may not hit their threshold, leaving their saving behavior unchanged. Another set of households may have been saving enough to send their children to college, $S \geq \tau$, prior to the education expansion. These high-saving households might have higher incomes or expected benefits from college, or they may have already faced a p that exceeded their threshold. High-saving households could even reduce their saving (but only as low as τ) as p goes up. We show this for the model in the Appendix (and for the data in Section 6.2), and the result is intuitive. An increase in p for a household already saving more than τ merely increases the chance of obtaining θ in period 2. This increase in expected income in period 2 lets the household consume more (and save less) in period 1, to smooth consumption. The increase in p may also reduce uncertainty, which again lowers saving.

We introduce the data next and then provide the details of our identification strategy.

4 The Data

We use household level information from the China Household Income Project (CHIP). 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.

[7] 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 raw data contains 6,929 household observations in 1995 and 6,835 in 2002. We focus on households with school-age children, 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 variable of interest is the saving rate, as calculated for each household using Equation (1). To eliminate outliers, we drop households in the tails (highest and lowest one percent) of the saving rate distribution in each year. We also drop the few households that changed Hukou after the expansion but before their children took the college entrance exam. Lastly, we 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 2,900 school-age children households and 1,218 college-age households in 1995 and 2,357 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. 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, their 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.[8] Many of the controls also had large changes over time. For

[8] Appendix B provides definitions for each variable.

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 (Chen et al., 2016; Berkowitz et al., 2017; Chen and 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.[9] Figure 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. 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 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 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. Our empirical approach overcomes the data limitations in two ways. First, we use the provincial enrollment rates in conjunction with the sample of households with college-age children to estimate the change in college probabilities for households with younger children. Second, we use a quasi 'difference-in-differences' regression approach to compare how the estimated changes in college probability affected saving

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

rates for households in 2002 (who were all impacted by the policy) versus how the probability changes affected 1995 households (before the education expansion occurred). Thus, we can net out unobserved factors that may cause a spurious relationship between saving rates and the increased likelihood of college, leaving only the impact attributable to the education expansion.

Specifically, we estimate Equation (6) twice; once using the 1995 sample of households with school-age children and once using the 2002 sample of households with school-age children.

$$SR_{ij} = \pi \Delta p_{ij} + X'_{ij}\delta + \lambda_j + u_{ij}. \tag{6}$$

The dependent variable is the saving rate (SR_{ij}) for household i from province j. 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 akin to a differences-in-differences estimate, with the difference attributable to the education expansion. Another interpretation is that the estimate of π using the 1995 households provides a 'counterfactual' for what would have happened to saving rates for the 2002 households in the absence of the education expansion.

The vector X_{ij} contains the control variables listed in Table 1, λ_j represents province fixed effects, and u_{ij} captures measurement error. Note that the coefficients (δ and λ's) are estimated separately for 1995 and 2002. The impact from the controls can vary over time. 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 this issue in Section 7.1.

We do not directly observe Δp_{ij}; we estimate it in two steps using the households with college-age children and the estimates of province enrollment rates described above. First, we estimate Equation (7) two times, once using households with college-age children in the 1995 sample and separately using the 2002 college-age households.

$$p_{ij} = X'_{ij}\beta + \gamma ER_i + \zeta_j + v_{ij}. \tag{7}$$

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. 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.

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 Equation (7) to get an estimate of p for each household with young children.[10] 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.

$$\Delta p_{ij} = X'_{ij} \cdot \left(\hat{\beta}^{2002} - \hat{\beta}^{1995}\right) + \left(\hat{\gamma}^{2002} ER_j^{2002} - \hat{\gamma}^{1995} ER_j^{1995}\right) + \left(\hat{\zeta}_j^{2002} - \hat{\zeta}_j^{1995}\right). \qquad (8)$$

We assume that the variation in Δp across households was exogenously driven by the higher education expansion (e.g. by the differences in province enrollment rates depicted in Figure 5). If this were strictly true, then we would only need to calculate Δp for the 2002 households and estimate Equation (6) 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 data to provide evidence that the saving rate trends were in fact similar across households with young children in the years leading up to the college expansion. Now, though, we turn to our main empirical results.

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

6 Higher Education Expansion's Effect on Saving

This section reports our main findings, which are obtained by estimating Equation (6) via ordinary least squares and Equation (7) as a probit model. According to our estimates, a 10 percentage point increase in the expected probability of attending college leads to a more than 7 percentage 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, we first estimate Equation (7) 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 Equation (7) to calculate an expected college probability for each household with young children (ages 6 to 18). Figure 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.

Finally, we calculate the change in college probability using Equation (8). 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). Figure 7 plots a fitted curve (smoothed by a local polynomial regression (Fan and 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 Equation (6). The estimate of π (the marginal effect of the change in college probability) for the 2002 sample equals 0.849. The effect on saving is quantitatively large and statistically different from 0 at better than the 1 percent significance level. Households for which the probability of attending college went up saved more. The estimate for the 1995 sample is much smaller (0.101) although still positive. 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 0.748, 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 saving rate by about 7.5 percentage points. This effect is very large, but given the dramatic increase in saving observed in China after 2000, it is not implausible.

Table 3 reports two sets of robustness checks. 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 Beijing and Chongqing (columns 3 and 4). 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 Figure 5 and Appendix Table C.2). 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.

Our measure of the saving rate could be partially determined by transitory shocks to household income (Carroll and Samwick, 1998). Following a strand of literature, in the second set of robustness checks (columns 5 and 6), we replace our dependent variable in Equation (6) with a measure of each household's cumulative saving rate.[11] The 1995 CHIP survey asks about annual income going back to 1990 and the 2002 survey asks about income every year starting with 1998. We aggregate each household's previous 4 years' incomes to form a measure of permanent income, in order to average

[11] See Lusardi (1998), Fuchs-Schündeln and Schuendeln (2005), and He et al. (2018), for example.

out transitory shocks. Then, the cumulative saving rate is defined as $\frac{W_i}{I_i}$, where W_i is household's total assets reported in the survey year and I_i is our measure of permanent income. With this set-up, the coefficient on Δp represents the impact on this accumulated saving rate. Again, we find that the increase in college probability leads households to save more. 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 (note, for columns 5 and 6, we drop households that do not report wealth), but the effects remain statistically significant and quantiatively large.

The empirical results in Tables 2 and 3 are consistent with the intuition coming out of our structural model that low-saving households ($S < \tau$) experiencing a large enough increase in the likelihood of college (p) will start to save more (i.e. Proposition 1). However, the model also implies that high-saving households could react to the college expansion by saving less (see Appendix A). In other words, the model suggests a difference in the effect on families with ample savings versus those with little. We can use the data along with our difference-in-differences approach to check this implication from the structural model.

To this end, we first regress gross savings (rather than saving rates) on the full set of observables and province fixed effects, separately for each year. We use the 1995 coefficient estimates to predict counterfactual saving levels for households in 2002 (and vice versa). Based on these predicted savings, we classify households into low savings (bottom quartile) and high savings (top quartile). Then, we re-run our main regressions for each group in each year.

Table 4 presents the results. The first column reports the estimates for the low saving and high saving households in 1995; the second column reports 2002. A positive correlation exists for low savings households in both years. In contrast, this relationship is negative for high savers in both years and large and statistically significant after the college expansion. These estimates are consistent with the idea (coming out of the structural model) that households with high saving levels reduce their saving when college opportunities increase. The DD estimates further indicate that an increased college probability leads high-saving households to save less, while low-saving households save more.[12]

We end this section by revisiting our identification strategy. A critical assumption for the

[12] Note, these results are consistent with our main findings, as we have omitted the middle two quartiles and Table 4 focuses on gross saving rather than saving rates.

difference-in-differences framework is that the differentially treated groups were not trending (before the treatment) in ways correlated with the treatment. In our case, the concern is that the types of households experiencing large increases in their college probability were also the types of households 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. Figure 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 (Equation (7)) were a function of the provincial enrollment rates. So, 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 the 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.[13] 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

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

remains large and statistically significant.[14]

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

$$SR_{ij} = W_{ij}'\delta + \pi \cdot \Delta p_{ij} + \phi \cdot Z_{ij} + \psi \cdot \Delta p_{ij} \cdot Z_{ij} + \lambda_j + u_{ij}, \tag{9}$$

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.[15] 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 5 provide evidence against this possibility. The coefficient estimate for Δp (0.786) in 2002 is high, even 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 (7.6 percentage points) remains large and close to our baseline result.

[14] 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 and Wooldridge, 2009).

[15] To save space, the tables below do not report the coefficient estimates for the controls.

7.1.2 Public Health Coverage

Healthcare is another motive for precautionary saving (Chamon and 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 5 report estimates controlling for heterogeneous effects based on having public health coverage. The resulting DD estimate of 0.89 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.

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 5 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 (0.80) 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 emobodied in Equation (9), 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 and 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 Equation (9), interacting the new variable ($Z = Single$) with the change in the probability of college.

Table 6 columns (1) and (2) present the results. The resulting DD estimate for the average marginal effect equals 0.77, close to our baseline result. Note, the coefficient estimate on *Single*, while not statistically significant, is close to 0.024 both before and after the education expansion. This estimate is similar to that in Lugauer et al. (2017); they estimate 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 prevelant 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.[16] Therefore, we next interact college probability with a variable ($Z = Male$) that equals one 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 6 suggest that the male child affect is quite small. Moreover, our estimate of the college expansion effect (DD=0.76) 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 ($Z = Age$) with the change in college probability. Columns (5) and (6) of Table 6 present the results. Once again, the DD estimate (0.88)

[16] Cai et al. (2019) and Zhou (2014) also discuss related issues.

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 characterisitics. 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. However, 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.

References

Bairoliya, Neha, Ray Miller, and Akshar Saxena (2018). The Macroeconomic Impact of Fertility Changes in an Aging Population. *Working paper*.

Banerjee, Abhijit, Esther Duflo, Rachel Glennerster, and Cynthia Kinnan (2015). The Miracle of Microfinance? Evidence from a Randomized Evaluation. *American Economic Journal: Applied Economics 7*(1), 22–53.

Becker, Gary S. and H. Gregg Lewis (1973). On the Interaction between the Quantity and Quality of Children. *Journal of Political Economy 81*(2, Part 2), S279–S288.

Berkowitz, Daniel, Hong Ma, and Shuichiro Nishioka (2017). Recasting the Iron Rice Bowl: The Reform of China's State-Owned Enterprises. *Review of Economics and Statistics 99*(4), 736–747.

Bollinger, Christopher and Chenxu Hu (2017). China's Higher Education Expansion and College Premium: IV Estimations. *Working paper*.

Bollinger, Christopher and Chenxu Hu (2018). Effects of Cohort Size on College Premium: Evidence from China's Higher Education Expansion. *Working paper*.

Brown, Meta, John Karl Scholz, and Ananth Seshadri (2012). A New Test of Borrowing Constraints for Education. *Review of Economic Studies 79*(2), 511–538.

Cai, Xiqian, Yi Lu, Jessica Pan, and Songfa Zhong (2019). Gender Gap Under Pressure: Evidence from China's National College Entrance Examination. *Review of Economics and Statistics 101*(2), 249–263.

Cameron, Stephen V. and Christopher Taber (2004). Estimation of Educational Borrowing Constraints Using Returns to Schooling. *Journal of Political Economy 112*(1), 132–182.

Carniero, Pedro and James J. Heckman (2009). The Evidence on Credit Constraints in Post-Secondary Schooling. *The Economic Journal 112*(482), 705–734.

Carroll, Christopher D. and Andrew A. Samwick (1998). How Important Is Precautionary Saving? *The Review of Economics and Statistics 80*(3), 410–419.

Chamon, Marcos D., Kai Liu, and Eswar S. Prasad (2013). Income Uncertainty and Household Savings in China. *Journal of Development Economics 105*, 164–177.

Chamon, Marcos D. and Eswar S. Prasad (2010). Why Are Saving Rates of Urban Households in China Rising? *American Economic Journal: Macroeconomics 2*(1), 93–130.

Chao, Chi Chur, Jean Pierre Laffargue, and Eden Yu (2011). The Chinese Saving Puzzle and the Life-Cycle Hypothesis: A Revaluation. *China Economic Review 22*(1), 108–120.

Che, Yi and Lei Zhang (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.

Chen, Binkai and Rudai Yang (2012). Understanding China's Rising Saving Rate: The Role of Higher Education Reform. In *The New Chinese Economy*, pp. 137–153. Springer.

Chen, Binkai, Xi Yang, and Ninghua Zhong (2016). Urban Housing Privatization and Household Saving in China. *Working paper*.

Chen, Kaiji and Yi Wen (2017). The Great Housing Boom of China. *American Economic Journal: Macroeconomics 9*(2), 73–114.

Choi, Horag, Steven Lugauer, and Nelson C. Mark (2017). Precautionary Saving of Chinese and U.S. Households. *Journal of Money, Credit and Banking 49*(4), 635–661.

Choukhmane, Taha, Nicholas Coeurdacier, and Keyu Jin (2017). The One-Child Policy and Household Savings. *Working paper*.

Curtis, Chadwick C. (2016). Economic Reforms and the Evolution of China's Total Factor Productivity. *Review of Economic Dynamics 21*, 225–245.

Curtis, Chadwick C., Steven Lugauer, and Nelson C. Mark (2015). Demographic Patterns and Household Saving in China. *American Economic Journal: Macroeconomics 7*(2), 58–94.

Curtis, Chadwick C., Steven Lugauer, and Nelson C. Mark (2017). Demographics and Aggregate Household Saving in Japan, China, and India. *Journal of Macroeconomics 51*, 175–191.

Doepke, Matthias and Michèle Tertilt (2016). Families in Macroeconomics. In *Handbook of Macroeconomics*, Volume 2. Elsevier.

Doepke, Matthias and Fabrizio Zilibotti (2019). *Love, Money, and Parenting: How Economics Explains the Way We Raise Our Kids.* Princeton University Press.

Fan, Elliott, Xin Meng, Zhichao Wei, and Guochang Zhao (2017). Rates of Return to Four-Year University Education: An Application of Regression Discontinuity Design. *The Scandinavian Journal of Economics*.

Fan, Jianqing and Irene Gijbels (1996). *Local Polynomial Modelling and Its Applications: Monographs on Statistics and Applied Probability*, Volume 66. CRC Press.

Feng, Shuaizhang and Xiaoyu Xia (2018, December). Endogenous Skill-Biased Technology Adoption: Evidence from China's College Enrollment Expansion Program. Working Papers 2018-099, Human Capital and Economic Opportunity Working Group.

Fuchs-Schündeln, Nicola and Matthias Schuendeln (2005). Precautionary Savings and Self-Selection: Evidence from the German Reunification "Experiment". *The Quarterly Journal of Economics 120*(3), 1085–1120.

Ge, Suqin, Dennis Tao Yang, and Junsen Zhang (2018). Population Policies, Demographic Structural Changes, and the Chinese Household Saving Puzzle. *European Economic Review 101*, 181–209.

Greenwood, Jeremy, Nezih Guner, and Guillaume Vandenbroucke (2017). Family Economics Writ Large. *Journal of Economic Literature 55*(4), 1346–1434.

He, Hui, Feng Huang, Zheng Liu, and Dongming Zhu (2018). Breaking the "Iron Rice Bowl": Evidence of Precautionary Savings from the Chinese State-Owned Enterprises Reform. *Journal of Monetary Economics 94*, 94–113.

Imbens, Guido W. and Jeffrey M. Wooldridge (2009). Recent Developments in the Econometrics of Program Evaluation. *Journal of Economic Literature 47*(1), 5–86.

Imrohoroglu, Ayse and Kai Zhao (2018a). Intergenerational Transfers and China's Social Security Reform. *Journal of the Economics of Ageing 11*, 62–70.

Imrohoroglu, Ayse and Kai Zhao (2018b). The Chinese Saving Rate: Long-Term Care Risks, Family Insurance, and Demographics. *Journal of Monetary Economics 96*, 33–52.

Keane, Michael P. and Kenneth I. Wolpin (2001). The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment. *International Economic Review 42*(4), 1051–1103.

Knight, John, Quheng Deng, and Shi Li (2017). China's Expansion of Higher Education: The Labour Market Consequences of a Supply Shock. *China Economic Review 43*, 127–141.

Li, Hongbin, Yueyuan Ma, Lingsheng Meng, Xue Qiao, and Xinzheng Shi (2017). Skill Complementarities and Returns to Higher Education: Evidence from College Enrollment Expansion in China. *China Economic Review 46*(March), 10–26.

Li, Shi and Chunbing Xing (2010). China's Higher Education Expansion and its Labor Market Consequences. *IZA Discussion Paper* (4974).

Lugauer, Steven and Nelson C. Mark (2013). The Role of Household Saving in the Economic Rise of China. *Global Interdependence, Decoupling, and Recoupling*, 131.

Lugauer, Steven, Jinlan Ni, and Zhichao Yin (2017). Chinese Household Saving and Dependent Children: Theory and Evidence. *China Economic Review*.

Lusardi, Annamaria (1998). On the Importance of The Precautinary Saving Motive. *The American Economic Review 88*(5), 449–453.

Malkova, Olga and Breno Braga (2018). Hope for the Family: The Effects of Merit Aid Programs on Parental Outcomes. *Working paper*.

Modigliani, Franco and Shi Larry Cao (2004). The Chinese Saving Puzzle and the Life-Cycle Hypothesis. *Journal of Economic Literature 42*(1), 145–170.

Rosenzweig, Mark R. and Junsen Zhang (2014). Co-residence, Life-Cycle Savings and Inter-Generational Support in Urban China. *NBER Working Paper* (20057).

Shen, Hong and Wenli Li (2003). *A Review of the Student Loans Scheme in China.*

Song, Zheng, Kjetil Storesletten, Yikai Wang, and Fabrizio Zilibotti (2015). Sharing High Growth across Generations: Pensions and Demographic Transition in China. *American Economic Journal: Macroeconomics 7*(2), 1–39.

Song, Zheng and Dennis T. Yang (2010). Life Cycle Earnings and the Household Saving Puzzle in a Fast-Growing Economy. *Manuscript, Chinese University of Hong Kong*.

Sun, Stephen Teng and Constantine Yannelis (2016). Credit Constraints and Demand for Higher Education : Evidence From Financial Deregulation. *Review of Economics and Statistics 98*(March), 12–24.

Wan, Yinmei (2006). Expansion of Chinese Higher Education Since 1998 : Its Causes and Outcomes. *7*(1), 19–31.

Wang, Xiaojun, Belton M. Fleisher, Haizheng Li, and Shi Li (2014). Access to College and Heterogeneous Returns to Education in China. *Economics of Education Review 42*, 78–92.

Wei, Shang-Jin and Xiaobo Zhang (2011). The Competitive Saving Motive: Evidence from Rising Sex Ratios and Savings Rates in China. *Journal of Political Economy 119*(3), 511–564.

Yao, Rui, Feifei Wang, Robert O. Weagley, and Li Liao (2011). Household Saving Motives: Comparing American and Chinese Consumers. *Family and Consumer Sciences Research Journal 40*(1), 28–44.

Yoo, Kyeongwon and John Giles (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, Weina (2014). Brothers, Household Financial Markets and Savings Rate in China. *Journal of Development Economics 111*, 34–47.

9 Tables

Table 1: Summary Statistics by Year and Children's Age

	School-Age		College-Age	
	1995	2002	1995	2002
Saving rate	0.107	0.142	0.141	0.128
	(0.226)	(0.342)	(0.238)	(0.334)
Child college attainment dummy			0.101	0.183
			(0.301)	(0.387)
Controls				
Number of kids	1.111	1.076	1.631	1.159
	(0.323)	(0.267)	(0.630)	(0.404)
Number of elderly	0.0838	0.0878	0.0755	0.0647
	(0.313)	(0.340)	(0.279)	(0.262)
Private house	0.410	0.782	0.475	0.757
	(0.492)	(0.413)	(0.500)	(0.429)
Housing accumulation fund	0.441	0.543	0.424	0.540
	(0.497)	(0.498)	(0.494)	(0.499)
Age	39.74	39.94	50.10	47.27
	(5.058)	(4.911)	(4.638)	(3.273)
Gender	0.633	0.652	0.664	0.638
	(0.482)	(0.477)	(0.472)	(0.481)
Currently employed	0.980	0.923	0.860	0.853
	(0.140)	(0.267)	(0.347)	(0.354)
College degree	0.066	0.110	0.100	0.050
	(0.247)	(0.313)	(0.300)	(0.219)
SOE job	0.822	0.325	0.825	0.342
	(0.383)	(0.468)	(0.380)	(0.475)
Public health	0.703	0.636	0.709	0.680
	(0.457)	(0.481)	(0.454)	(0.467)
Employment tenure	15.80	13.58	21.41	16.32
	(7.425)	(8.613)	(10.23)	(11.60)
Years of schooling	10.45	11.46	10.32	10.29
	(2.984)	(3.004)	(3.651)	(2.881)
Spouse years of schooling	9.657	10.67	8.714	9.582
	(3.476)	(3.571)	(4.148)	(3.439)
Annual income	1.349	2.244	1.655	2.335
	(0.705)	(1.408)	(0.900)	(1.394)
Annual expenses	1.175	1.805	1.379	1.901
	(0.624)	(1.222)	(0.797)	(1.133)
Total assets	1.076	3.954	1.454	4.056
	(1.473)	(10.05)	(3.932)	(5.415)
Observations	2900	2357	1218	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.

Table 2: College Probability's Effect on Household Saving Rates

	(1) 1995	(2) 2002
Δp	0.101	0.849
	(0.099)	(0.240)
No. kids	0.009	0.016
	(0.009)	(0.018)
No. elderly people	0.003	0.009
	(0.008)	(0.006)
Private house	0.000	-0.042
	(0.009)	(0.016)
Housing accumulation fund	0.013	0.012
	(0.006)	(0.010)
Head age	0.001	-0.003
	(0.001)	(0.001)
Head male	0.001	0.026
	(0.009)	(0.010)
Head currently working	0.020	0.040
	(0.014)	(0.022)
Head with college degree	-0.002	-0.057
	(0.015)	(0.021)
Head SOE job	0.019	0.019
	(0.006)	(0.012)
Head public health	0.002	0.030
	(0.009)	(0.009)
Head tenure	0.000	0.000
	(0.000)	(0.001)
Head years of schooling	-0.000	0.001
	(0.001)	(0.001)
Spouse years of schooling	0.004	0.007
	(0.001)	(0.002)
Annual income	0.433	0.277
	(0.062)	(0.026)
Annual expenses	-0.531	-0.359
	(0.080)	(0.033)
Total assets	0.001	-0.004
	(0.002)	(0.001)
Observations	2900	2357
R^2	0.745	0.698
DD	0.748	
	(0.260)	

Notes: This table reports regression estimates of Equation 6 for 1995 and 2002. Row DD reports the difference in the coefficient estimates (π) on Δp. The regressions include province fixed effects, and the parentheses report robust standard errors clustered by province.

Table 3: Robustness Checks

	No Beijing		No Beijing & Chongqing		Cumulative Saving	
	(1)	(2)	(3)	(4)	(5)	(6)
	1995	2002	1995	2002	1995	2002
Δp	0.133	0.614	0.149	0.675	1.453	9.895
	(0.063)	(0.215)	(0.067)	(0.258)	(0.468)	(5.265)
Observations	2690	2225	2690	2136	2536	2232
R^2	0.735	0.693	0.735	0.689	0.085	0.263
DD	0.481		0.526		8.442	
	(0.224)		(0.267)		(5.286)	

Notes: Each regression is based on Equation 6. The dependent variable for the first four columns remains the saving rate. The dependent variable for columns 5 and 6 is the cumulative saving rate. Each regression includes the full set of controls, except assets are not included in columns 5 and 6 due to collinearity with the dependent variable. The regressions include province fixed effects, and the parentheses report robust standard errors clustered by province.

Table 4: Impact on Low and High Savers

	(1) 1995	(2) 2002	(3) DD
Low	0.129	0.547	0.418
	(0.304)	(0.718)	(0.780)
	[725]	[590]	
High	-1.015	-8.280	-7.265
	(0.821)	(1.902)	(2.071)
	[725]	[589]	

Notes: This table reports the effect on high (predicted savings above the 75th percentile) and low (below the 25th percentile) saving households using gross savings as the dependent variable. Each regression includes province fixed effects and the full set of controls. Robust standard errors are in parentheses and clustered by province. Brackets contain the number of observations.

Table 5: Other Reforms

	SOE job		Public health		Private house	
	(1)	(2)	(3)	(4)	(5)	(6)
	1995	2002	1995	2002	1995	2002
ΔP	0.022	0.786	-0.096	1.057	0.126	1.092
	(0.174)	(0.238)	(0.078)	(0.492)	(0.104)	(0.301)
$\Delta P \times$SOE job	0.076	0.187				
	(0.176)	(0.159)				
$\Delta P \times$ Public health			0.201	-0.193		
			(0.035)	(0.269)		
$\Delta P \times$ Private house					-0.073	-0.248
					(0.075)	(0.194)
SOE job	0.018	0.011	0.018	0.019	0.019	0.019
	(0.006)	(0.015)	(0.005)	(0.012)	(0.006)	(0.012)
Public health	0.002	0.031	0.001	0.036	0.002	0.031
	(0.009)	(0.009)	(0.009)	(0.011)	(0.009)	(0.009)
Private house	0.000	-0.042	0.002	-0.045	0.000	-0.040
	(0.009)	(0.016)	(0.009)	(0.019)	(0.009)	(0.016)
Observations	2900	2357	2900	2357	2900	2357
R^2	0.745	0.698	0.745	0.698	0.745	0.698
Marginal effect	0.085	0.846	0.046	0.934	0.096	0.897
	(0.094)	(0.237)	(0.092)	(0.334)	(0.099)	(0.242)
DD	0.761		0.888		0.801	
	(0.255)		(0.346)		(0.261)	

Notes: Each regression includes province fixed effects and the control variables. The parentheses report robust standard errors clustered by province.

Table 6: Demographics

	Number of kids		Gender of kids		Life-cycle	
	(1)	(2)	(3)	(4)	(5)	(6)
	1995	2002	1995	2002	1995	2002
Δp	0.075	1.118	0.038	0.858	-0.061	2.122
	(0.102)	(0.356)	(0.142)	(0.241)	(0.466)	(0.909)
Single	0.024	0.025				
	(0.049)	(0.037)				
$\Delta p\times$ Single	0.028	-0.269				
	(0.116)	(0.237)				
Male			0.004	0.003		
			(0.003)	(0.006)		
$\Delta p\times$ Male			0.105	-0.017		
			(0.099)	(0.121)		
Head age					0.001	-0.002
					(0.001)	(0.001)
$\Delta p\times$ Head age					0.004	-0.029
					(0.009)	(0.019)
Observation	2900	2357	2900	2357	2900	2357
R^2	0.745	0.698	0.745	0.698	0.745	0.699
Marginal effect	0.100	0.869	0.089	0.850	0.079	0.957
	(0.096)	(0.242)	(.109)	(0.239)	(0.143)	(0.259)
DD	0.769		0.762		0.877	
	(0.260)		(0.263)		(0.296)	

Notes: Each regression includes province fixed effects and the control variables. The parentheses report robust standard errors clustered by province.

10 Figures

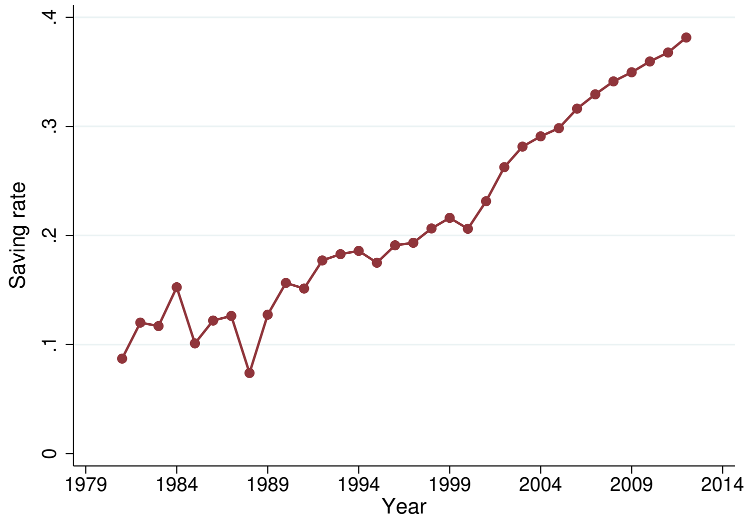

Figure 1: Household Saving Rates

Notes: The data is from the China Yearly Statistical Book and the National Bureau of Statistics.

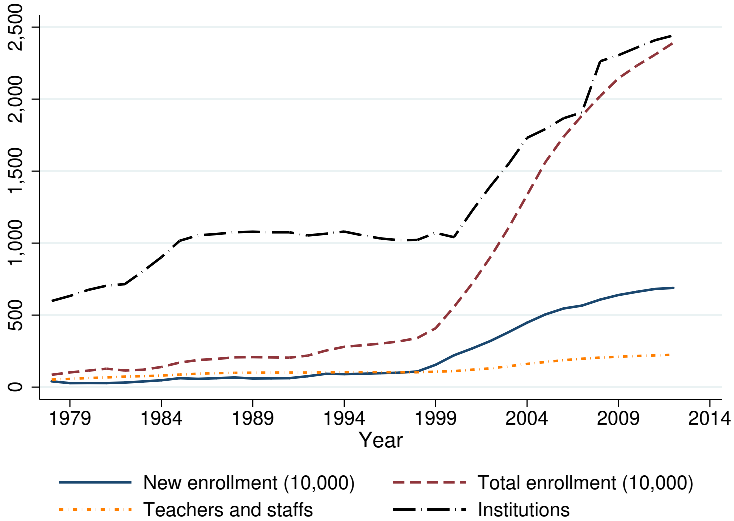

Figure 2: China Higher Education Expansion

Notes: The data comes from the China Yearly Statistical Book and the National Bureau of Statistics.

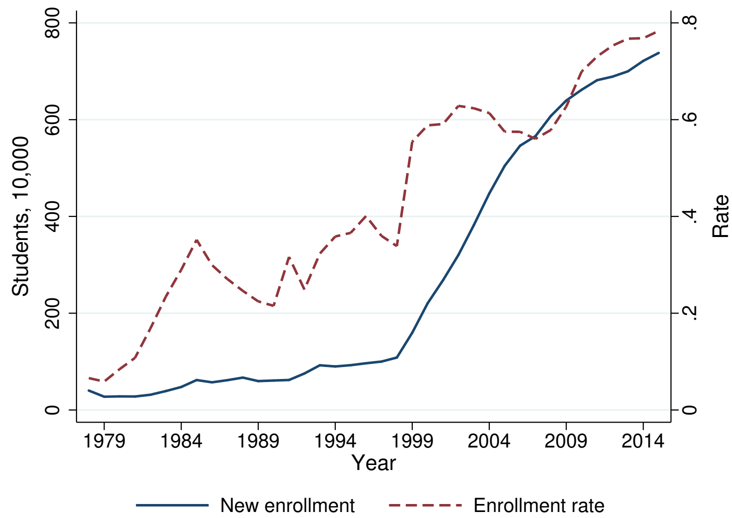

Figure 3: New Students Enrolled and the Enrollment Rate

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.

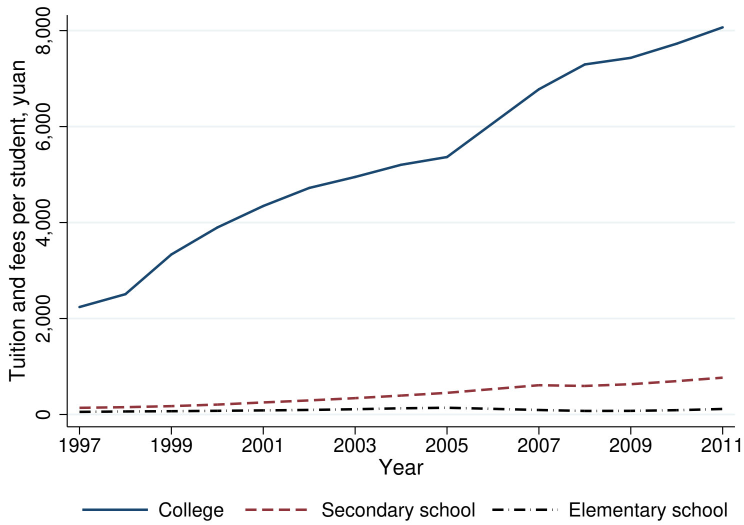

Figure 4: Cost of Education

Notes: The data comes from the China Yearly Statistical Book and the National Bureau of Statistics.

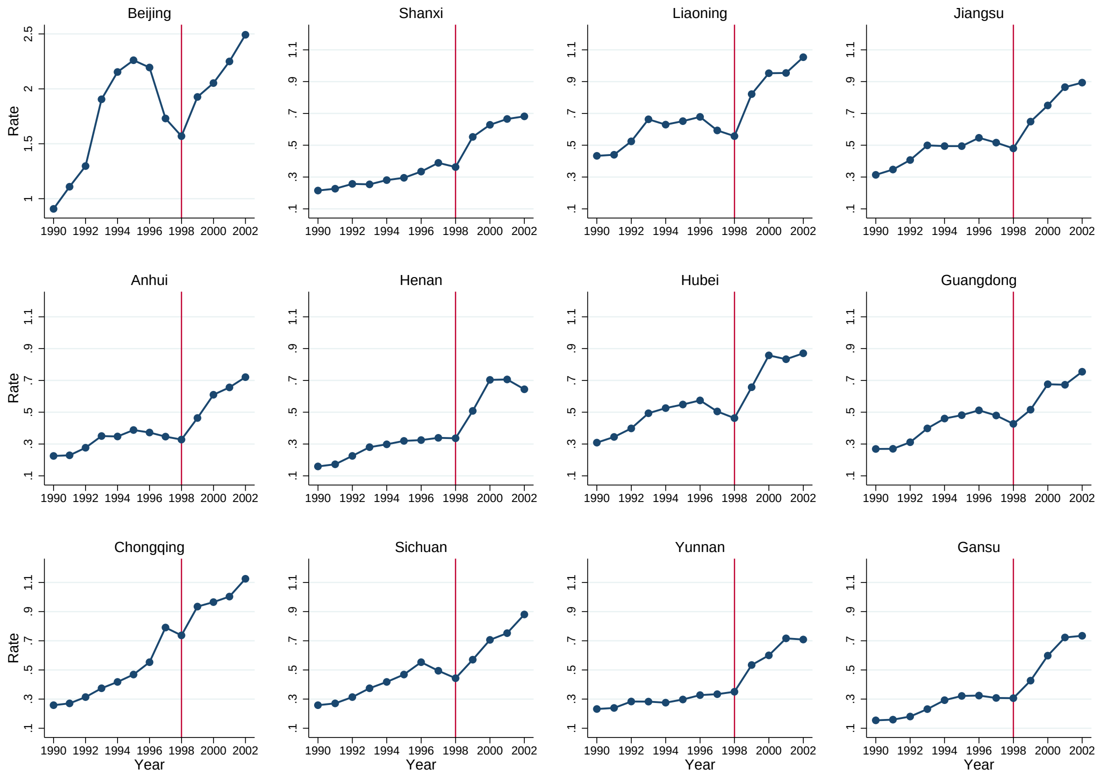

Figure 5: Enrollment Rate by Province

Note: 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.

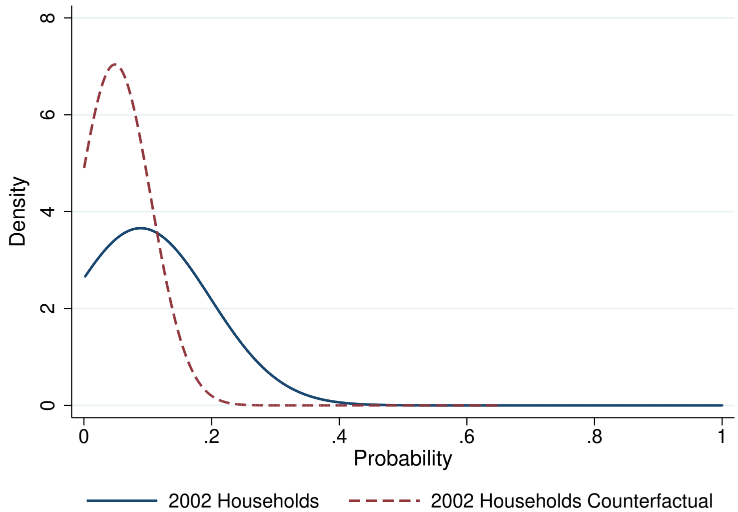

Figure 6: Density of College Probability for 2002 Households

Note: 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 Equation 7.

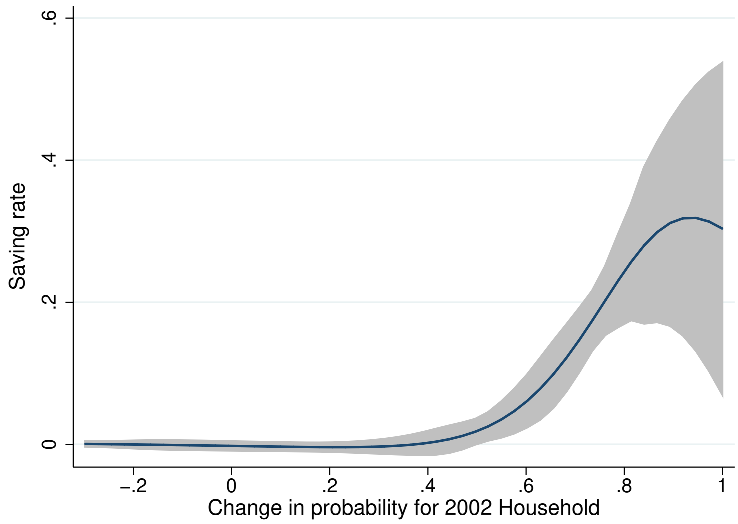

Figure 7: Saving Rate and Change of Probability for 2002 Households

Note: This figure shows the relationship between saving rates and the change in college probability for 2002 households via local polynomial regressions.

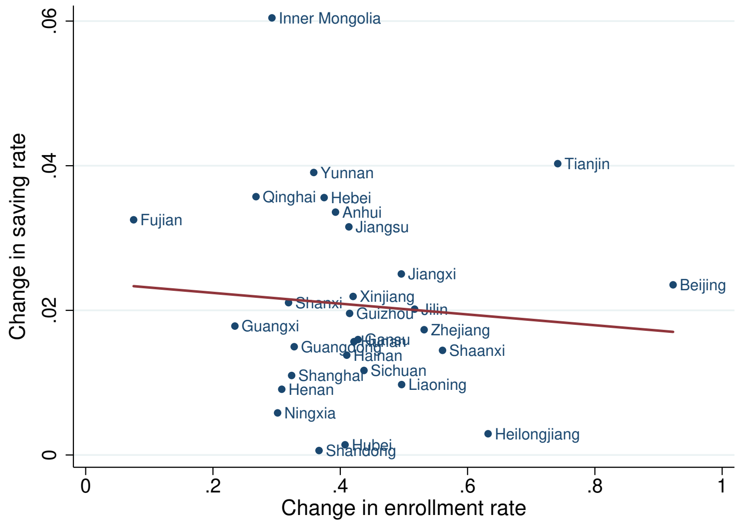

Figure 8: Test for Pre-Trends

Note: 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.