

The One-Child Policy and Household Saving*

Taha Choukhmane
Yale

Nicolas Coeurdacier
SciencesPo and CEPR

Keyu Jin
London School of Economics

This Version: July 7, 2017

Abstract

We investigate whether the ‘one-child policy’ has contributed to the rise in China’s household saving rate and human capital in recent decades. In a life-cycle model with intergenerational transfers and human capital accumulation, fertility restrictions lower expected old-age support coming from children—inducing parents to raise saving and education investment in their offspring. Quantitatively, the policy can account for at least 30% of the rise in aggregate saving. Using the birth of twins under the policy as an empirical out-of-sample check to the theory, we find that quantitative estimates on saving and education decisions line up well with micro-data.

Keywords : Life Cycle saving, Fertility, Human Capital, Intergenerational Transfers.

JEL codes: E21, D10, D91

*We thank Pierre-Olivier Gourinchas, Nancy Qian, Andrew Chesher, Aleh Tsyvinski, and seminar participants at LSE, SciencesPo, HEI Geneva, Cambridge University, SED (Seoul), CREI, Banque de France, Bilkent University, University of Edinburgh, EIEF, IIES, Yale, UC Berkeley for helpful comments. Taha Choukhmane: 27 Hillhouse Avenue, 06510, New Haven, CT, USA; email: taha.choukhmane@yale.edu; Nicolas Coeurdacier thanks the ERC for financial support (ERC Starting Grant INFINHET) and the SciencesPo-LSE Mobility Scheme. Contact address: SciencesPo, 28 rue des saint-pères, 75007 Paris, France. email: nicolas.coeurdacier@sciencespo.fr ; Keyu Jin: London School of Economics, Houghton Street, WC2A 2AE, London, UK; email: k.jin@lse.ac.uk.

1 Introduction

The one-child policy, introduced in 1979 in urban China, was one of the most radical birth control schemes implemented in history. The policy, aimed at curbing the high population growth, limited each urban household to one child. The consequence was a drastic decline in the urban fertility rate over a short period of time—from on average 3 children per family in the late 1960s to just about 1 in the early 1980s. The radical implementation of the one-child policy made it a natural experiment in Chinese history, albeit to date an under-studied event.

In this paper, we examine the quantitative effects of the one-child policy on Chinese saving and human capital – building up from its micro-level impact at the household level to its aggregate implications. China’s household saving rate has been increasing at a rapid rate: between 1982 and 2014, the average urban household saving rate rose steadily from 12% to 31%. Human capital accumulation has also accelerated over the last thirty years, with the average years of schooling increasing by about 50%—from 5.8 years to 8.9 for an adult aged 25 (Barro and Lee (2010); see also Li et al. (2013)).

In the Chinese society, children act as a source of old-age support. Parents rear and educate children when young, while children make financial transfers and provide in-kind benefits to their retired parents. Not only is the custom commonplace, it is also stipulated by constitutional law. How many children one decides to have directly affects the amount of transfers parents receive. Imagine that families that typically had 3 children were suddenly constrained to 1. The reduction in expected transfers means that parents now have to save more on their own. In other words, parents shift their investment in the form of children towards the form of financial assets. This is what we call the ‘transfer channel’.

Additionally, the reduction in overall expenditures owing to fewer children also raises the household saving rate. When education costs can amount to 5 to 15% of household income per child depending on its age, the fall in expenditures from having fewer children can be substantial. These additional resources are partly saved—what we label as the ‘expenditure channel’. Both channels tend to exert upward pressure on the household saving rate and constitute the micro-channels of the policy on saving. On the aggregate level, demographic compositional changes associated with a fall in fertility rates also affect the aggregate saving rate—as is well-understood through the classic formulations of the life-cycle motives for saving (Modigliani (1986)). Our approach shows that the aforementioned micro-channels on saving are more important in the Chinese context—where inter-generational transfers within families are large in magnitude.

The second consequence is that the one-child policy may have led to a rapid accumulation of human capital of the only child generation. When parents can substitute quantity for quality, the expected reduction in transfers implied by the policy can be partly compensated by raising the child’s education investment and expected future income. The importance of the interaction between saving and human capital decisions is thus immediately apparent: the degree of substitution of quantity for quality determines the impact on saving of the one-child policy. In other words, if parents can perfectly compensate for quantity with quality—say, if human capital adjusts at no cost—then the policy would have little effect on saving, and the transfer channel, in particular, would disappear.

In investigating the joint impact of the one-child policy on human capital and saving, the paper makes three main contributions: (i) providing a tractable model linking fertility, intergenerational transfers and human capital accumulation; (ii) expanding it to a quantitative framework that can be calibrated to micro data; (iii) conducting an empirical test of the theory using the births of twins as exogenous deviations from the policy.

Specifically, the theoretical framework incorporates two new elements to the standard lifecycle theory of saving: intra-family transfers and human capital accumulation. Agents make decisions on the number of children to bear, the level of human capital to endow them, and on how much to save for retirement. Children are costly, but at the same time, presents an investment opportunity by offering support to their parents at a later stage. An exogenous reduction in fertility lowers total expenditures spent on children and raises household saving ('expenditure channel'); this holds notwithstanding a substitution of 'quantity' for 'quality'—with more education spending on the only child. The rise in the child's future wages owing to human capital accumulation is in general not enough to compensate for the overall reduction in transfers that parents receive when retired, providing further incentives to save ('transfer channel').

Our model thus sheds light on the interaction between human capital and saving decisions. A stronger policy response of human capital—driven for instance by weaker diminishing returns to education—severely limits the saving response. Also, we show that under certain conditions, one can identify the micro-channel on saving and the human capital response over time through a *cross-sectional* comparison of twin households and only-child households. This forms the basis of our later empirical analysis and counterfactual exercises.

Our second contribution lies in the quantitative investigation of our theory. The model is expanded and calibrated to micro-level Chinese data. Starting from aggregate implications, we find that the model imputes at least a 30% and at most 60% of the rise in the household saving rate over 1982-2014 to the one-child policy—depending on the natural fertility rate that would have prevailed without the policy change. Matching predicted human capital accumulation to the data is less straightforward, though our model predicts that the policy has significantly increased the human capital of the only child generation by at least 24% compared to their parents.

Our multi-period model implies different saving behavior across age groups. Taking one step further, we examine the evolution in the *age-saving profile* over time. We find that our model can capture quantitatively the overall shift in saving rates across ages. We also show that the evolution of the profile is, however, vastly inconsistent with the predictions from a standard OLG model without old-age support and human capital accumulation. In the absence of the transfer channel, saving of parents in their 50s (whose children have departed from households) should have fallen following the policy—the opposite of what is observed in the data.

Finally, the predictions of the model at the micro-level are evaluated through a 'twin experiment', which serves as an 'out-of-sample' test to the quantitative performance of the model. In this experiment, we compare the cross-sectional differences in saving and education spending between only-child and twin families with the differences estimated from micro-data. Using the births of twins as an exogenous fertility shock is appealing under the one-child policy since households must have one child and randomly, sometimes, they have two (twins). Our empirical results reveal that twin households save on average 5 to 8 percentage points less (as a % of income) than only-child households. This difference remains once children have left the household, indicating that the transfer channel is at play. While education expenditures (as a % of income) are about 6 percentage points higher in twin households, education expenditures *per child* are about 2 percentage points less on twins than on an only child—with twins being less educated. Overall, the proximity of the empirical findings to model estimates suggests reasonable quantitative predictability of our model.

Related literature. Our paper closely relates to the literature explaining the staggeringly high saving rate in China, starting with Modigliani and Cao (2004) ('Chinese Saving Puzzle'). In a

sense, a distinguishing feature of our paper is our endeavor to bridge the micro-level approach with the macro-level approach.¹ The ability to match the micro-evidence gives further credence to the model’s macroeconomic implications. Storesletten and Zilibotti (2013) provide an exposition of the transformation of the Chinese society and the perplexingly high household saving in the recent years, and discusses some recent developments in the literature.² Our paper relates to theoretical work linking fertility and saving starting with Barro and Becker (1989),³ but also focuses on the interaction between human capital and saving decisions. The interaction is quantitatively critical for our results and largely absent in those studies.⁴ Note also that the nature of intergenerational altruism differs from that of Barro and Becker (1989)—in our view, the assumption that parents rear children to provide for old-age more aptly captures the family arrangements of a developing country like China than the notion that children’s lives are a continuation of their parents’. Finally, our paper builds on a large literature linking fertility changes and human capital accumulation, from theory (starting with Becker and Lewis (1973)) to the use of twin births as identification strategy (Rosenzweig and Wolpin (1980)).⁵ Our theory, however, differs from the quantity-quality trade-off derived from utility assumptions, as it appears endogenously in the presence of old-age support.

A few caveats are in order. The form of intergenerational transfers occurs within households in this economy, in contrast to intergenerational transfers taking place through social security—which has until now been virtually non-existent in China, and unreliable to say the least. We treat these transfers towards the elderly as a social norm and thus exogenously given in our model, contrary to Boldrin and Jones (2002). Our model also treats interest rates as exogenous and abstracts from general equilibrium effects of saving on capital accumulation and interest rates. We believe this to be realistic in the Chinese context where households face interest rates largely determined by the government.⁶ A theoretical general equilibrium analysis may be found in Banerjee et al. (2014) and our subsequent work (Coeurdacier et al. (2014)).

The paper is organized as follows. Section 2 provides certain background information and facts that motivate some key assumptions underlying our framework. Section 3 provides our theoretical model that links fertility, education and saving decisions in an overlapping generations model. Section 4 develops a calibrated quantitative model to simulate the impact of the policy. The empirical tests based on twins and model counterfactuals are conducted in Section 5. Section 6 concludes.

¹Modigliani and Cao (2004), Horioka and Wan (2007), Curtis, Lugauer, and Mark (2015) find some evidence supporting the link between demographics and saving at the aggregate level, but meet difficulty when confronting micro-data. Focusing on long-term care risk, a recent paper by Imrohoroglu and Zhao (2017) goes further in inspecting the transfer channel through which fertility affects saving. They also provide comforting micro-evidence.

²Some compelling explanations of the saving puzzle include: (1) precautionary saving (Blanchard and Giavazzi (2005), Chamon and Prasad (2010) and Wen (2011)); (2) habit formation (Carroll and Weil (1994)); (3) changes in income profiles (Song and Yang (2010), Guo and Perri (2012)); (4) gender imbalances and competition in the marriage market (Wei and Zhang (2011) and Du and Wei (2013)); (5) demographics (Modigliani and Cao (2004), Horioka and Wan (2007), Curtis, Lugauer, and Mark (2015), Banerjee et al. (2014) and Imrohoroglu and Zhao (2017)); (6) income growth and credit constraints (Coeurdacier, Guibaud and Jin (2015)), interacted with housing costs (Bussiere et al. (2013) and Wan (2015))—though there is no consensus as to whether the rise in housing prices can explain the rise in household saving (Wang and Wen (2012)). Chamon and Prasad (2010) and Yang, Zhang and Zhou (2011) provide a thorough treatment of facts pertaining to China’s saving, and at the same time present the challenges that some of these theories face.

³See also Boldrin and Jones (2002), Chakrabarti (1999), Cigno and Rosati (1996), Raut and Srinivasan (1994).

⁴Manuelli and Seshadri (2007) extend Barro and Becker (1989) to include human capital but do not explore the role of saving.

⁵See Angrist et al. (2010) for references. Rosenzweig and Zhang (2009) also use the birth of Chinese twins to measure the ‘quantity-quality’ trade-off in children; they find supporting evidence of our model’s predictions (see also Hongbin et al. (2008), Oliveira (2012) and Qian (2013)).

⁶Despite capital controls, China is also a semi-open economy where household saving is largely channeled abroad.

2 Motivation and Background

Based on various aggregate and household level data sources from China, this section provides stylized facts on (1) the background of the ‘one-child policy’ and its consequences on the Chinese demographic composition; (2) the direction and magnitude of intergenerational transfers—from parents to children in financing their education, and from children to parents in support of their old age. The quantitative relevance of these factors motivates the main assumptions underlying the theoretical framework. Micro and macro data sources used are described in Appendix A.

2.1 The One-Child Policy and the Chinese demographic transition

The one-child policy decreed in 1979 was intended to curb the high population growth in the Maoist China of the 1950s-1960s. The consequence was a sharp drop in the nation-wide fertility rate. The policy was strictly enforced in urban areas and partially implemented in rural provinces.⁷ Binding fertility constraints is a clear imperative for the purpose of our study and urban households are therefore a natural focal point in our analysis. It is important to note that the rise in saving in China is mostly driven by urban households, which account for 88% of the increase between 1982-2014.⁸

The one-child policy and the demographic evolution in the 1970s. Starting from 1971, the Chinese government promoted family planning to reduce population growth. These initiatives were captured by the slogan ‘wan, xi, shao’ (later, longer, fewer) that encouraged postponing marriage until a later age, lengthening birth spacing between children, and reducing their number (Cai (2010) and Scharping (2003)). The timing and the extent of enforcement of these policies varied across regions and a significant discretion was given to local governments to implement them. In the late 1970s, the Chinese government shifted to a stricter approach of population planning imposing a limit on the number of children per couple: a two children limit implemented nationwide in 1978 (Scharping (2003)) followed by the one-child policy announced in 1979 and strictly enforced in urban areas after 1980. As shown in Figure 1 (upper-panel), in a span of three years, the share of first-birth in total births jumped from a fairly stable share of 55% in 1977 to 90% in 1981, while the share of higher-order births declined symmetrically.

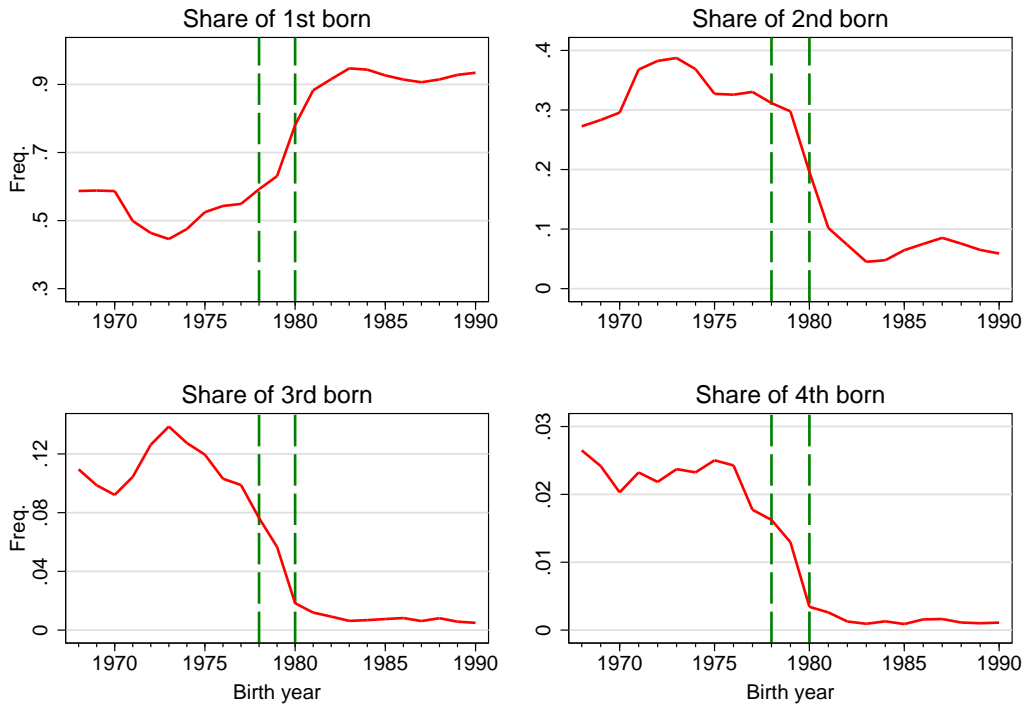
Due to this large shock to fertility behavior between 1978 and 1980, the completed fertility by date of birth of children fell from roughly three in 1970 to about one ten years later (Figure 1, bottom-panel). At this point, it is crucial to understand that the child limits imposed in the late 1970s also affected household who started to conceive earlier on—explaining the progressive decline shown in Figure 1 (bottom-panel). Indeed, parents having their first child in the 1970s, before the policy, were also constrained in their ability to have additional children later on. The reason is that it takes time to conceive multiple children. For instance, a couple with a first child born in 1975 would conceive a second one, on average, 3 years later. By the time they would likely conceive a third child, the one-child policy would have kicked in, reducing their completed fertility. Applying this reasoning for every household with a first-born in the 1970s, we show in Appendix B that the one-child policy can account for the gradual decrease in fertility for parents who had children in the

⁷Household-level data (Urban Household Survey, UHS) confirm a strict enforcement of the One Child Policy for urban households: over the period 2000-2009, 96% of urban households that had children had only one child. Some urban households had more than one child. If we abstract from the birth of twins, accounting for about 1% of households, the remaining 3% of households may include minority ethnicities (not subject to the policy)—accounting for a sufficiently small portion to be discarded.


⁸Urban household saving rate grew by about 20 percentage points over the period, whereas rural household saving rate barely changed. Source: CEIC.

Figure 1: The one-child policy and fertility in urban China

The 1978-80 fertility shock



Fertility by average date of birth of children



Notes: The upper-panel shows the number of births of a n -th child divided by the total number of births in a given year. The vertical lines corresponds to a two children limit in 1978 and the one-child policy in 1980. The bottom-panel shows the completed fertility by average date of birth of children. At a given date t , it shows the number of children in households whose average date of birth of children is equal to t . The number of children only includes surviving children. Data source: Census, restricted sample where only urban households are considered. See Appendix A.

1970s.⁹ Additional evidence of the major role played by the policy in constraining fertility is provided in the same Appendix when comparing the fertility of the Han (main ethnic group) and the non-Han (minority) populations. While both groups had similar fertility in 1970, the non-Hans had one more child in the 1980s as they were only subject to a two children limit (Figure B.6 in Appendix B). This strongly suggests that policies limiting the number of children, either to one or two, are crucial in explaining the fertility behavior of Chinese urban families.

The demographic structure since 1980. The demographic structure evolved accordingly, ensuing fertility controls (Table 1). Some prominent patterns are: (1) a sharp rise in the median age— from 22 years in 1980 to 37 years in 2015; (2) a rapid decline in the share of young individuals (ages 0-19) from 47% to 23% over the period, and (3) a corresponding increase in the share of middle-aged population (ages 30-59). While the share of the young is expected to drop further until 2050, the share of the older population (above 60) increases sharply only after 2015— when the generation of the only-child ages. In other words, the one-child policy leads first to a sharp fall in the share of young individuals relative to middle-aged adults, followed by a sharp increase in the share of the elderly only one generation later.

Table 1: Demographic structure in China

	1980	2015	2050
Share of young (age 0-19/Total Population)	47%	23%	19%
Share of middle-aged (age 30-59/Total Population)	28%	45%	36%
Share of elderly (age above 60/Total Population)	8%	15%	35%
Median age	22	37	48

Note: Data source: UN World Population Prospects (2017).


2.2 Intergenerational Transfers

Old-age support. Intergenerational transfers from children to elderly are the bedrock of the Chinese society. Beyond cultural norms, it is also stipulated by Constitutional law: “children who have come to age have the duty to support and assist their parents” (Article 49). Failure in this responsibility may even result in law suits. According to Census data in 2005, family support is the main source of income for almost half of the elderly (65+) urban population (Figure 2, left panel). From the China Health and Retirement Longitudinal Study (CHARLS), individuals of ages 45-65 in 2011 expect this pattern to continue in the coming years: half expect transfers from their children to constitute the main source of income for old age (Figure 2, right panel).

CHARLS provides further detailed data on intergenerational transfers in 2008 for two provinces: Zhejiang (a prosperous coastal province) and Gansu (a poor inland province). We restrict the sample to urban households in which at least one member (respondent or spouse) is older than 60 years of age. Old age support takes broadly two forms: financial transfers (‘direct’ transfers) and ‘indirect’

⁹Assuming that all parents had the same fertility and birth spacing behaviors as those with a first born in 1964 (thus presumably barely affected by fertility policies), our counterfactual exercise presented in Appendix B documents that the 1978-1980 policies can, alone, account for nearly all of the fall in fertility of parents with a first birth in the 1970s. See Figure B.3 in Appendix B. Appendix B also provides evidence that the early seventies ‘wan, xi, shao’ policy had a quantitatively small impact on fertility.

Figure 2: Main Source of Livelihood for the Elderly (65+) in urban areas



Notes: Left panel, Census (2005). Right panel, CHARLS (2011), urban households, whole sample of adults between 45-65 (answer to the question: Whom do you think you can rely on for old-age support?).

Table 2: Transfers towards elderly: Descriptive Statistics

Number of households	321
Average number of adult children (25+)	3.5
Share living with adult children	44%
Incidence of positive net transfers	
- from adult children to parents	77%
- from parents to adult children	4%
Net transfers in % of parent's total income	
- All parents	51%
- Transfer receivers only	61%
Of which households with:	
- One or two children	16%
- Three children	46%
- Four children	68%
- Above Five children	80%


Notes: Data source: CHARLS (2008). Restricted sample of urban households with a respondent/spouse of at least 60 years of age with at least one surviving adult children aged 25 or older. Transfers is defined as the sum of regular and non-regular financial transfers in yuan. Net Transfers are transfers from children to parents less the transfers received by children. Parent's total income is defined as the sum of positive net transfers received from children plus income from employment, pensions and asset returns.

transfers in the form of co-residence or other in-kind benefits. According to Table 2, 44% of the elderly reside with their children in urban households. Positive (net) transfers from adult children to parents occur in 77% of households and are large in magnitude—constituting the largest share of old-age income of on average 51% of all elderly's income (and up to 61% if one focuses on the sample of transfer receivers). Table 2 also shows that transfers (as a % of total income) are increasing in the number of children. The flip side of the story is that restrictions in fertility will therefore likely reduce the amount of transfers conferred to the elderly. This fact bears the central assumption underlying our theoretical framework.

Education expenditures. An important feature of our theory is that education expenditures for children are important for understanding saving across age-groups and over time, following fertility

changes. Education expenses are a prominent source of transfers from parents towards their children according to the Chinese Household Income Project (CHIP) in 2002.¹⁰ Restricting our attention to families with an only child, Figure 3 displays education expenditures (in % of household income) in relation to the age of the child; it increases from roughly 5% for a child below 10 up to 10-15% for a child above 13. Data provides some evidence on the relative importance of ‘compulsory’ and ‘non-compulsory’ (or discretionary) education costs: not surprisingly, the bulk of expenditures (about 80%) incurred for children above 16 can be considered as discretionary, whereas the opposite holds for younger children.¹¹ This evidence motivates the assumption that education costs are more akin to a compulsory cost (per child) for young children, while it is more of a choice variable subject to a quantity-quality trade-off for older children.

Figure 3: Education Expenditures for a child, by age of the child (% of household income)




Notes: Data source: CHIP (2002). Sample restricted to urban households with an only child. This graph plots the average expenditure (as a share of household income) across education categories by the age of the only child.

Timing of transfers from children to parents. The timing and direction of transfers—paid and received at various ages of adulthood (computed from CHARLS (2008))—guide the assumptions adopted by the quantitative model. Figure 4 (left panel) displays the evolution of the average net transfers of children to parents (in monetary values; left axis) as a function of the (average) age of children. The right panel displays the net transfers received by parents as a function of their age. Observing the left panel, one can mark that net transfers are on average negative at young ages (children receiving transfers from parents), and increase sharply at the age of 25. This pattern accords

¹⁰The Urban Household Surveys (UHS) in 2006 and RUMiCI in 2008 show a similar pattern although estimates are slightly larger in magnitude.

¹¹Compulsory education costs are mostly kindergarten/nursery, tuition and fees for compulsory education, textbooks. Discretionary costs include mostly non-compulsory education tuition and fees. See Appendix A for details.

Figure 4: Timing of intergenerational transfers



Notes: CHARLS (2008), whole sample of urban households. The left panel plots the average amount of net transfers of children to his/her parents (left axis) and the % of coresidence (right axis) by the average age of child. The right panel plots the average amount of net transfers received by parents from their children (left axis) and the % of coresidence (right axis) by the average age of parents.

with the notion that education investment is the main form of transfers towards children. After this age, children confer increasing amounts of transfers towards their parents—received by parents upon retirement (right panel). If co-residence (right axis) is also considered as a form of transfers, a similar pattern emerges: children leave the parental household upon reaching adulthood (left panel).¹² For parents in their 60s, the degree of co-residence no longer falls with parental age, remaining around 40–50% as parents return to live with their children (right panel).

3 Theoretical Analysis

We develop a tractable multi-period overlapping generations model with intergenerational transfers, endogenous fertility and human capital accumulation. The parsimonious model yields a semi-closed form solution that serves two main purposes. First, it reveals the fundamental channels driving the fertility-human capital-saving relationships. Second, the model motivates our empirical strategy, showing how one can identify the impact of the one-child policy on human capital accumulation and saving through a cross-sectional comparison between two-children (twin) households and only-child households. A quantitative version of the model is developed in the subsequent section, although the main mechanisms are elucidated in the following model.

¹²Co-residence is the focus of Rosenzweig and Zhang (2014), which analyzes to what extent the young people’s option of co-residing with their parents affect saving decisions.

3.1 Set-up

Consider an overlapping generations economy in which agents live for four periods, characterized by: childhood, youth (y), middle-age (m), and old-age (o).

Timing. An individual born in period $t - 1$ does not make decisions on his consumption in childhood, which is assumed to be proportional to parental income. The agent supplies inelastically one unit of labor in youth and in middle-age, and earns a wage rate $w_{y,t}$ and $w_{m,t+1}$, which is used, in each period, for consumption, transfers and asset accumulation $a_{y,t}$ and $a_{m,t+1}$. At the end of period t , the young agent makes the decision on the number of children n_t to bear and on the amount of human capital h_t to endow each of his children. In middle-age, in $t + 1$, he transfers a combined amount of $T_{m,t+1}$ to his n_t children and parents—to augment human capital of the former, and consumption of the latter. In old-age, the agent consumes all available resources, coming from gross returns on accumulated assets $a_{m,t+1}$ and transfers from children $T_{o,t+2}$.

Preferences and budget constraints. An individual maximizes the life-time utility which includes the consumption $c_{\gamma,t}$ at each age γ and the benefits from having n_t children:

$$U_t = \log(c_{y,t}) + v \log(n_t) + \beta \log(c_{m,t+1}) + \beta^2 \log(c_{o,t+2})$$

where $v > 0$ reflects the preference for children, and $0 < \beta < 1$. The sequence of budget constraints for an agent born in $t - 1$ obeys

$$\begin{aligned} c_{y,t} + a_{y,t} &= w_{y,t} \\ c_{m,t+1} + a_{m,t+1} &= w_{m,t+1} + Ra_{y,t} - T_{m,t+1} \\ c_{o,t+2} &= Ra_{m,t+1} + T_{o,t+2}. \end{aligned} \tag{1}$$

Agents lend (or borrow) through bank deposits, earning a constant and exogenously given gross interest rate R .¹³ Because of parental investment in education, the individual born in period $t - 1$ enters the labor market with an endowment of human capital h_{t-1} . Assuming decreasing returns parametrized by $0 < \alpha < 1$, the human capital h_{t-1} , along with an experience parameter $e < 1$, and a deterministic level of economy-wide productivity z_t , determines the wage rates:

$$\begin{aligned} w_{y,t} &= ez_t h_{t-1}^\alpha \\ w_{m,t+1} &= z_{t+1} h_{t-1}^\alpha. \end{aligned} \tag{2}$$

Intergenerational transfers. The cost of raising kids is assumed to be paid by parents in middle-age, in period $t + 1$, for a child born at the end of period t . The total cost of raising n_t children is proportional to current wages, $n_t \phi(h_t) w_{m,t+1}$, where $\phi(h) = \phi_0 + \phi_h h$, $\phi_0 > 0$ and $\phi_h > 0$. The ‘mouth to feed’ cost, including consumption and compulsory education expenditures (per child), is a fraction ϕ_0 of the parents’ wage rate; the discretionary education cost $\phi_h h_t$ is increasing in the level of human capital chosen by the parents.

Transfers made to the middle-aged agent’s parents amount to a fraction $\psi n_{t-1}^{\omega-1} / \omega$ of current wages $w_{m,t+1}$, with $\psi > 0$ and $0 < \omega \leq 1$. This fraction is decreasing in the number of siblings—to capture

¹³This is analogous to a model in which the central bank intermediates household saving abroad. This modelling choice is adopted for the purpose of distilling the most essential forces governing the fertility-saving relationship without undue complication of the model. This is also reasonable in the Chinese context, where interest rates on households deposits are largely set by the government.

the possibility of free-riding among siblings sharing the burden of transfers. We treat these transfers as an institutional norm in China; children supporting their parents is not only socially expected, but is even stipulated by law. The assumed functional form for transfers is analytically convenient, but (i) its main properties are tightly linked to the data (see Section 4.2); (ii) these properties are also qualitatively retained with endogenous transfers but at the expense of tractability and facility of parametrization.¹⁴

The combined amount of transfers made by the middle-aged agent in period $t + 1$ to his children and parents thus satisfy: $T_{m,t+1} = (n_t \phi(h_t) + \psi n_{t-1}^{\omega-1} / \omega) w_{m,t+1}$. An old-age parent receives transfers from his n_t children: $T_{o,t+2} = \psi \frac{n_t^\omega}{\omega} w_{m,t+2}$.

3.2 Household decisions and model dynamics

Consumption decisions. Optimal consumption can be solved given fertility and human capital decisions. The following assumption,

Assumption 1 *The young are subject to a credit constraint, binding in all periods:*

$$a_{y,t} = -\theta \frac{w_{m,t+1}}{R}$$

specifies that the young can borrow up to a constant fraction θ of the present value of future wage income. For a given θ , the constraint is more likely to bind if productivity growth is high (relative to R) and the experience parameter e is low. This assumption is necessary for obtaining a realistic saving behavior of the young—one that avoids a counterfactual sharp borrowing that emerges under fast growth and a steep income profile (see also Coeurdacier, Guibaud and Jin (2015)).

Assumption 1 and the absence of bequests mean that the only individuals that optimize their saving are the middle-aged.¹⁵ The assumption of log utility implies that the optimal consumption of the middle-age is a constant fraction of the present value of lifetime resources, which consist of current disposable income—net of debt repayments and current transfers to children and parents—and the present value of transfers to be received in old-age:

$$c_{m,t+1} = \frac{1}{1 + \beta} \left[\left(1 - \theta - n_t \phi(h_t) - \psi \frac{n_{t-1}^{\omega-1}}{\omega} \right) w_{m,t+1} + \frac{\psi}{R} \frac{n_t^\omega}{\omega} w_{m,t+2} \right]. \quad (3)$$

It follows from Eq. 1 that the optimal asset holding of a middle-aged individual is

$$a_{m,t+1} = \frac{\beta}{1 + \beta} \left[\left(1 - \theta - n_t \phi(h_t) - \psi \frac{n_{t-1}^{\omega-1}}{\omega} \right) w_{m,t+1} - \frac{\psi}{\beta R} \frac{n_t^\omega}{\omega} w_{m,t+2} \right]. \quad (4)$$

Eq. 4 illuminates the link between fertility and saving: parents with more children accumulate less wealth because they have less available resources for saving (term $n_t \phi(h_t)$) and because they expect

¹⁴In the data, transfers given by each child are indeed decreasing in the number of offspring, and the income elasticity of transfers is close to 1—as is assumed by the transfer function (see Section 4.2). In a model in which transfers are endogenously determined—where children place a weight on parents' old-age utility of consumption—the main properties hold in the steady-state: transfers are decreasing in the number of offspring, and the income elasticity of transfers is 1. While parents may desire to save less knowing that more saving beget less transfers from children, this effect amounts to a reduced discount rate. See also Boldrin and Jones (2002) for a model with endogenous old-age support.

¹⁵Without data on bequests in China, a bequest motive would be difficult to calibrate on top of existing parameters in our quantitative version of the theory developed in Section 4. Note that Horioka (2014) finds a significantly weaker bequest motive in China than in the US.

larger transfers (last term).

Fertility and Human Capital. Fertility decisions hinge on equating the marginal utility of bearing an additional child with the net marginal cost of raising the child:

$$\begin{aligned}\frac{v}{n_t} &= \frac{\beta}{c_{m,t+1}} \left(\phi(h_t) w_{m,t+1} - \frac{\psi n_t^{\omega-1} w_{m,t+2}}{R} \right) \\ &= \frac{\beta}{c_{m,t+1}} \left(\phi(h_t) - \mu_{t+1} \psi n_t^{\omega-1} \left(\frac{h_t}{h_{t-1}} \right)^\alpha \right) w_{m,t+1},\end{aligned}\quad (5)$$

where $\mu_{t+1} \equiv z_{t+2}/Rz_{t+1} \equiv (1 + g_{z,t+1})/R$ is the productivity growth-interest rate ratio. The right hand side is the net cost, in utility terms, of having an additional child. The net cost is the current marginal cost of rearing a child, $\partial T_{m,t+1}/\partial n_t$ less the present value of the benefit from receiving transfers next period from an additional child, $\partial T_{o,t+2}/\partial n_t$. In this context, children are analogous to investment goods—and incentives to procreate depend on the factor μ_{t+1} —productivity growth relative to the gross interest rate. Higher productivity growth raises the number of children—by raising future benefits relative to current costs. But saving in assets is an alternative form of investment, which earns a gross rate of return R . Thus, the decision to have children as an investment opportunity depends on this *relative* return.¹⁶

The optimal choice on the children's endowment of human capital h_t is determined by

$$\frac{\psi n_t^\omega}{R \omega} \frac{\partial w_{m,t+2}}{\partial h_t} = \phi_h n_t w_{m,t+1},$$

where the (discounted) marginal gain of having children more educated and thus providing more old-age support is equalized to the marginal cost of further educating them. Using Eq. 2, the above expression yields the optimal choice for h_t , given n_t and the predetermined parent's own human capital h_{t-1} :

$$h_t = \left[\frac{\psi}{\omega \phi_h} \frac{\alpha \mu_{t+1}}{h_{t-1}^\alpha n_t^{1-\omega}} \right]^{\frac{1}{1-\alpha}}. \quad (6)$$

A greater number of children n_t reduces the gains from educating them—a quantity and quality trade-off. This trade-off arises from the fact that the marginal benefit in terms of transfers is decreasing in the number of children ($\omega < 1$). Given any number of children n_t , incentives to provide further education is increasing in the productivity growth relative to the interest rate μ_{t+1} —which gauges the relative benefits of investing in children. Greater altruism ψ of children for parents also increases parental investment in them.

The optimal number of children n_t , combining Eq. 3, 5 and 6, satisfies, with $\lambda = \frac{v+\omega\beta(1+\beta)}{\alpha v+\alpha\beta(1+\beta)}$:

$$n_t = \left(\frac{v}{\beta(1+\beta) + v} \right) \left(\frac{1 - \theta - \psi \frac{n_{t-1}^{\omega-1}}{\omega}}{\phi_0 + \phi_h (1 - \lambda) h_t} \right). \quad (7)$$

Equations 6 and 7 are two equations that describe the evolution of the two state variables of the economy $\{n_t; h_t\}$. Eq. 6 describes the human capital response to a change in fertility n_t —with h_t decreasing in n_t . Eq. 7 measures the response of fertility to a change in the children's human capital h_t . There are two competing effects governing this relationship: the first effect is that higher levels of

¹⁶All else constant, the relationship between fertility and interest rates is negative—as children are considered as investment goods. This relationship is the opposite of the positive relationship in a dynastic model (Barro and Becker (1989)).

education per child raises transfers per child, motivating parents to have more children. The second effect is that greater education, on the other hand, raises the cost per child, and reduces the incentives to have more children. The first effect dominates if diminishing returns to transfers are relatively weak compared to diminishing returns to education, $\lambda > 1$ —in which case n_t is increasing in h_t .


Steady-State. The steady state is characterized by a constant productivity growth-interest rate ratio, $\mu_t = \mu$, and constant state variables $h_t = h_{ss}$ and $n_t = n_{ss}$. Eqs. 6 and 7 are, in the long run:

$$\frac{n_{ss}}{1 - \theta - \psi n_{ss}^{\omega-1} / \omega} = \left(\frac{v}{\beta(1 + \beta) + v} \right) \left(\frac{1}{\phi_0 + \phi_h (1 - \lambda) h_{ss}} \right) \quad (\text{NN})$$

$$h_{ss} = \left(\frac{\psi \alpha \mu}{\phi_h} \right) \frac{n_{ss}^{\omega-1}}{\omega}. \quad (\text{QQ})$$

Figure 5 depicts graphically the two curves for an illustrative calibration. The (NN) curve describes the response of fertility to higher education. Its positive slope (for $\lambda > 1$) captures the greater incentive of bearing children when they have higher levels of human capital. The downward sloping curve (QQ) shows the combination of n and h that satisfies the quantity/quality trade-off in children.

Figure 5: Steady-State Human Capital and Fertility Determination



Notes: Steady-state, with an illustrative calibration using $\phi_0 = 0.1$, $\phi_h = 0.1$, $\psi = 0.2$, $\beta = 0.986$ (per annum, 0.75 over 20 years), $R = 4\%$ (per annum), $g_z = 4\%$ (per annum), $\theta = 0$, $\omega = 0.7$, $\alpha = 0.4$. $v = 0.055$ set such that $n_{ss} = 3/2$.

Assumption 2 Parameters are restricted such that $\omega \geq \alpha$, implying $\lambda > 1$.

Assumption 2 ensures model convergence to a stable steady-state—avoiding divergent dynamics whereby parents constantly reduce their children’s education for cost reduction and increase their number (or vice-versa). This leads to the following proposition:

Proposition 1 There is a unique steady-state for the number of children $n_{ss} > 0$ and their human capital $h_{ss} > 0$ to which the dynamic model defined by Eqs. 6 and 7 converges. Also, comparative statics yield

$$\frac{\partial n_{ss}}{\partial \mu} > 0 \text{ and } \frac{\partial h_{ss}}{\partial \mu} > 0; \frac{\partial n_{ss}}{\partial v} > 0 \text{ and } \frac{\partial h_{ss}}{\partial v} < 0; \frac{\partial n_{ss}}{\partial \phi_0} < 0 \text{ and } \frac{\partial h_{ss}}{\partial \phi_0} > 0.$$

Proof: See Appendix C.

Higher productivity growth relative to the interest rate increases the incentives to invest in children, both in terms of quantity and quality. A stronger preference towards children (or lower costs of raising them) makes parents willing to have more children, albeit less educated (lower ‘quality’) ones.

3.3 The One-Child Policy

Fertility constraint. The government is assumed to enforce a law that compels each agent to have up to a number n_{\max} of children over a certain period $[t_0; t_0 + T]$ with $T \geq 1$. In the case of the one-child policy, the maximum number of children per individual is $n_{\max} = 1/2$. We now examine the transitory dynamics of the key variables following the implementation of the policy, starting from an initial *steady-state* of unconstrained fertility characterized by $\{n_{t_0-1}; h_{t_0-1}\}$, with $n_{t_0-1} > n_{\max}$. The additional constraint $n_t \leq n_{\max}$ is now added to the original individual optimization problem. We focus on the interesting scenario in which the constraint is binding ($n_t = n_{\max}$ for $t_0 \leq t \leq t_0 + T$). Under constrained fertility, one needs an additional assumption for the model to converge if $T \rightarrow \infty$:

Assumption 3 $\alpha < 1/2$.

Assumption 3 is necessary to avoid divergent paths of human capital accumulation where higher education increases expected transfers and gives further incentives to raise education without any offsetting feedback on fertility decisions. Note that the assumed values for α are well within the range of the macro literature (Mankiw et al. (1992) and survey by Sianesi and van Reenen (2000)).

3.3.1 Human Capital and Aggregate saving

Human capital. The policy aimed at reducing the population inadvertently *increases* the level of per-capita human capital, thus moving the long-run equilibrium along the (QQ) curve, as shown in Figure 5 and stated by the following Lemma:

Lemma 1 As $T \rightarrow \infty$, human capital converges to a new (constrained) steady-state h_{\max} such that:

$$h_{\max} = \left(\frac{\psi \alpha \mu}{\phi_h} \right) \frac{n_{\max}^{\omega-1}}{\omega} > h_{t_0-1}.$$

The first generation of only child also features higher level of human capital than their parents:

$$\frac{h_{t_0}}{h_{t_0-1}} = \left(\frac{n_{t_0-1}}{n_{\max}} \right)^{\frac{1-\omega}{1-\alpha}} > 1.$$

Proof: See Appendix C.

Aggregate saving. The aggregate saving of the economy is the sum of the aggregate saving of each generation $\gamma = \{y, m, o\}$ coexisting in a given period t . The aggregate saving to aggregate labour income ratio defines the aggregate saving rate s_t — a weighted average of the young, middle-aged and

old's individual saving rates, where the weights depend on both the population and relative income of the contemporaneous generations (see Appendix C for details). Assuming constant productivity growth to interest rate ratio μ , the impact of the one-child policy on the dynamics of the aggregate saving rate between t_0 and $t_0 + 1$ is given by the following Proposition:

Proposition 2 With binding fertility constraints in period t_0 , the aggregate saving rate increases *unambiguously* over a generation:

$$s_{t_0+1} - s_{t_0} > 0.$$

Proof: See Appendix C.

For a given level of human capital of the generation of only child h_{t_0} , the change in aggregate saving rate over the period after the implementation of the policy can be written as,

$$\begin{aligned} s_{t_0+1} - s_{t_0} = & \underbrace{\frac{(n_{t_0-1} - n_{\max})e}{1 + n_{\max}e} s_{t_0} + \frac{1}{1 + n_{\max}e} \theta \mu \left(n_{t_0-1} - n_{\max} \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha \right)}_{\text{macro-channel (composition effects)}} \quad (8) \\ & + \underbrace{\frac{1}{1 + n_{\max}e} \frac{\beta}{1 + \beta} \left[\phi_0 (n_{t_0-1} - n_{\max}) + \left(\alpha + \frac{1}{\beta} \right) \frac{\psi \mu}{\omega} \left(n_{t_0-1}^\omega - n_{\max}^\omega \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha \right) \right]}_{\text{micro-channel}}. \end{aligned}$$

where the initial steady-state aggregate saving rate s_{t_0} is given in Appendix C. The expression can be decomposed into a macro-channel and a micro-channel. The macro-economic channels comprise changes in the composition of population, and the composition of income attributed to each generation. A fall in fertility of size $(n_{t_0-1} - n_{\max})$ reduces the proportion of young borrowers, relative to the middle-aged savers (population composition); it also places more weight on the aggregate income attributed to the middle-aged savers of the economy and less to young borrowers (income composition), although the latter effect depends on the endogenous human capital response h_{t_0} . In our framework, the response of human capital does not offset the fall in fertility for $\omega > \alpha$ such that both forces exert upward pressure on the aggregate saving rate (see Appendix C for a proof).¹⁷

The micro-channel corresponds to the change in saving of middle aged-parents and encapsulates two effects. The first effect is the reduction in the total cost of children— fewer ‘mouths to feed’ (the first term $\phi_0 (n_{t_0-1} - n_{\max})$) and a fall in total (discretionary) education costs— in spite of the rise in human capital per child (the second term multiplied by ‘ α ’). The second effect is the ‘transfer channel’, and captures the need to save more with a reduction in expected old-age support —again, despite higher human capital per child (the third term multiplied by ‘ $1/\beta$ ’). Indeed, incorporating the response of human capital h_{t_0} , we get:

$$n_{t_0-1}^\omega - n_{\max}^\omega \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha = n_{t_0-1}^\omega \left(1 - \left(\frac{n_{\max}}{n_{t_0-1}} \right)^{\frac{\omega-\alpha}{1-\alpha}} \right) \geq 0$$

The response of human capital does not offset the fall in fertility such that total discretionary education expenditures and expected transfers fall with fewer children, leading to an unambiguous rise in middle-aged saving.¹⁸ However, the size of the response of human capital of only child is essential to assess

¹⁷In period $t_0 + 1$, the reduction in fertility has not yet fed into an increase in the proportion of the dependent elderly (relative to the middle-aged). Thus, the negative effect of the rising share of the elderly on the aggregate saving rate materializes only once the generation of only child reaches middle-age (at $t_0 + 2$).

¹⁸On top of the rising share of elderly to middle-aged, another effect absent during the transition is the lowering of middle-aged saving due to the fact that there are fewer siblings among whom the burden of supporting parents can be shared. However, this effect only shows up when the only child generation turns middle age and does not apply to middle-aged

quantitatively the response of aggregate saving. With a stronger response of human capital ($\alpha \rightarrow \omega$), the transfer channel disappears and the fall in expenditures is limited to the ‘mouths to feed’ term. To the opposite, with constant (exogenous) human capital, one might overstate the response of saving as shown in Eq. 8.

3.3.2 Identification Through ‘Twins’

We next show theoretically how one can identify the microeconomic channel (over time) through a *cross-sectional* comparison between only-child households and twin-households. Proofs of these results are relegated to Appendix C. Consider the scenario in which some middle-aged individuals exogenously deviate from the one-child policy by having twins. Two main testable implications regarding human capital and saving can be derived.

Quantity-Quality Trade-Off. Parents of twins devote less resources for education *per-child* but their *overall* discretionary education expenditures are higher:

$$\frac{1}{2} \leq \left(\frac{h_{t_0}^{twin}}{h_{t_0}} \right) = \left(\frac{1}{2} \right)^{\frac{1-\omega}{1-\alpha}} < 1. \quad (9)$$

The quantity-quality trade-off driving human capital accumulation can be identified by comparing twins and an only-child. This ratio as measured by the data also provides some guidance on the relative strength of ω and α . Despite the trade-off, the fall in human capital per capita is less than the increase in the number of children, so that total discretionary education costs are higher for twins (and are the same when $\alpha \rightarrow \omega$).

Identifying the micro-channel on saving. The micro-economic impact of having twins on the middle-age parent’s saving rate comprise the same ‘expenditure channel’ and ‘transfer channel’. Parents of twins save less and the difference in the saving rate between parents of an only-child and parents of twins in $t_0 + 1$ satisfies:

$$s_{m,t_0+1} - s_{m,t_0+1}^{twin} = \frac{\beta}{1+\beta} \left[n_{max} \phi_0 + \left(\alpha + \frac{1}{\beta} \right) \frac{\psi \mu}{\omega} n_{max}^\omega \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha \left(2^{\frac{\omega-\alpha}{1-\alpha}} - 1 \right) \right] > 0.$$

A Lower Bound for the Micro-Channel. Let $\Delta s_m = s_{m,t_0+1} - s_{m,t_0}$, the policy implied change in the saving rate of middle-aged parents, one generation after the policy implementation (second-term above bracket in Eq. 8). Δs_m reflects the micro-economic impact on saving of moving from unconstrained fertility n_{t_0-1} to n_{max} . One can estimate the micro-channel of the policy by comparing, in the cross-section, the saving behavior of parents of twins versus parents of only child:

Lemma 2 If the fertility rate in absence of fertility controls is two children per household ($n_{t_0-1} = 2n_{max}$), then

$$\Delta s_m = s_{m,t_0+1} - s_{m,t_0+1}^{twin}.$$

Proof: See Appendix C.

If the unconstrained fertility is 2 children per household, we can identify precisely the micro-economic impact of the policy—by comparing the saving rate of a middle-aged individual with an only child to the one of parents having twins. We can also deduce a *lower-bound* estimate for the overall impact of the policy on the saving rate of the middle-aged—if the unconstrained fertility is greater than 2 (as

parents in $t_0 + 1$.

in China prior to the policy change). That is, if $n_{t_0-1} > 2n_{\max}$, then

$$\Delta s_m > s_{m,t_0+1} - s_{m,t_0+1}^{twin}.$$

These theoretical results demonstrate that cross-sectional observations from twin-households can inform us of the impact of the one-child policy on saving behavior over time.

3.4 Discussion

Before turning to the quantitative implications of our theory, we discuss two potential caveats of our framework.

Identification. The identification strategy based on twins coming out of our model relies on a set of important assumptions: having two children that are expected or having twins leads to identical saving and education decisions; and, if some households can avoid the policy by manipulating fertility (having twins), and these households make different saving and education decisions compared to the average, then any empirical strategy based on twins would be biased. The validity of these assumptions is discussed in the empirical Section 5. Also, our theory shows how cross-sectional observations from twin-households is informative about the time-series change in saving following the policy. Strictly speaking, this result holds in our model if the natural fertility rate had not changed from prior to the policy. But as income in China has been rising rapidly, fertility most likely would have fallen even without the one-child policy—albeit at a slower speed. We discuss the potential evolution of fertility in the absence of policies in the context of our quantitative model in Section 5.2.

Partial equilibrium. Our theory assumes an exogenous real interest rate. Due to financial repression in China, most of the wealth of households is held in the form of deposits, with interest rates controlled by the government and kept artificially low (see Allen et al. (2015) and Song et al. (2011, 2015)). While the Chinese institutional environment justifies this approach, our theory neglects general equilibrium effects through which fertility changes could affect the interest rate and in turn modify saving decisions. Such general equilibrium effects, emphasized in Banerjee et al. (2014), could potentially mitigate the impact of fertility on saving. In our quantitative model of Section 4, we investigate the relevance of our assumption in the Chinese context using measures of the real rate faced by households.

4 A Quantitative OLG Model

We develop a multi-period quantitative version of our theory, calibrated to household-level data. A reasonably parameterized model can assess the quantitative impact of the one-child policy on aggregate saving and human capital over the period 1982-2014. In addition, it is able to deliver finer predictions of saving rates over the life-cycle and provide directly testable evidence that motivates our empirical Section 5.

4.1 Set-up and model dynamics

Timing. Agents live for γ_d periods, so that γ_d age-groups $\gamma = \{1, 2, \dots, \gamma_d\}$ coexist in the economy in each period. The timing of the events that take place over the lifecycle is similar to before: the agent is a child for the first $\underline{\gamma} - 1$ periods and starts working at age $\underline{\gamma}$. He makes fertility and human

capital decisions for his children at age $\gamma_n \geq \underline{\gamma}$. After giving birth to children, and before age $\bar{\gamma}$, he is rearing and educating children while making transfers to his elderly parents. He reaches old age at age $\bar{\gamma}$, with $\gamma_n < \bar{\gamma} \leq \gamma_d$ — age at which he starts receiving transfers from his children. In old age, he finances consumption from the previous saving and from the support of his children, dying with certainty at the end of period γ_d without leaving any bequests.¹⁹

Preferences. Let $c_{\gamma,t}^i$ denote the consumption of an individual aged γ in period t , with $\gamma \in \{\underline{\gamma}, \underline{\gamma} + 1, \dots, \gamma_d\}$. The lifetime utility of an agent born at t entering the labor market at date $t + \underline{\gamma}$ is

$$U(t) = v \log(n_{t+\gamma_n}) + \sum_{\gamma=\underline{\gamma}}^{\gamma_d} \beta^{\gamma-\underline{\gamma}} \log(c_{\gamma,t+\gamma}), \quad (10)$$

with $0 < \beta < 1$ and $v > 0$. $n_{t+\gamma_n}$ denotes the number of children the agent has at date $t + \gamma_n$.

Life income profile and transfers. An individual born at t and entering the labor market at date $t + \underline{\gamma}$ with human capital H_t earns $w_{\gamma,t+\gamma} = e_{\gamma} z_{t+\gamma} H_t^{\alpha}$ at age γ and date $t + \gamma$. His human capital depends on the level of his parents $H_{t-\gamma_n}$, and their human capital investment h_t : $H_t = h_t^{1-\rho} H_{t-\gamma_n}^{\rho}$ with $\rho \in [0; 1]$ measuring the intergenerational transmission of human capital — $\rho = 0$ in the model of Section 3. e_{γ} is an experience factor of the life income profile; $z_{t+\gamma}$ represents aggregate productivity and is assumed to be growing at a constant rate of $z_{t+1}/z_t = 1 + g_z$.

The functional form of transfers and the costs of rearing and educating children are retained from before, although the timing of expenditures is more elaborate. Data reveals the timing and scale of these expenditures and transfers. We assume education costs are paid from age γ_n until age $\gamma_n + \gamma_e$. For an agent born at date t , children's compulsory education costs paid at age $\gamma \in \{\gamma_n, \dots, \gamma_n + \gamma_e\}$ are a fraction $\phi_{\gamma} n_{t+\gamma_n}$ of the agent's wage income $w_{\gamma,t+\gamma}$. The discretionary education costs are borne at the same age and are a fraction $\phi_{\gamma,h} h_{t+\gamma_n} n_{t+\gamma_n}$ of the wage income — $h_{t+\gamma_n}$ denotes the investment in human capital decided by the parents of the children born at date $t + \gamma_n$.

Transfers to support parents are made at age $\gamma \in \{\bar{\gamma} - \gamma_n, \dots, \gamma_d - \gamma_n\}$ and are a fraction $\psi \frac{n_t^{\omega-1}}{\omega}$ of the wage income. When old, at age $\gamma \geq \bar{\gamma}$, the agent receives transfers from his $n_{t+\gamma_n}$ children equal to $\psi \frac{n_{t+\gamma_n}^{\omega}}{\omega} w_{\gamma-\gamma_n,t+\gamma}$.

We denote $T_{\gamma,t+\gamma}$ the net transfers paid at age γ and date $t + \gamma$, which is the sum of transfers made to children and parents net of transfers received from children in old age:

$$T_{\gamma,t+\gamma} = \left[\mathbf{1}_{\{\gamma_n \leq \gamma \leq \gamma_n + \gamma_e\}} (\phi_{\gamma} + \phi_{\gamma,h} h_{t+\gamma_n}) n_{t+\gamma_n} + \mathbf{1}_{\{\bar{\gamma} - \gamma_n \leq \gamma \leq \gamma_d - \gamma_n\}} \psi \frac{n_t^{\omega-1}}{\omega} \right] w_{\gamma,t+\gamma} - \mathbf{1}_{\{\bar{\gamma} \leq \gamma \leq \gamma_d\}} \psi \frac{n_{t+\gamma_n}^{\omega}}{\omega} w_{\gamma-\gamma_n,t+\gamma}$$

where $\mathbf{1}_{\{x \leq \gamma \leq y\}}$ is equal to one if $\gamma \in \{x, \dots, y\}$ and zero otherwise.

Budget and credit constraints. An agent born at date t and of age γ faces the following instantaneous budget constraint at each age γ :

$$a_{\gamma,t+\gamma} = w_{\gamma,t+\gamma} - c_{\gamma,t+\gamma} - T_{\gamma,t+\gamma} + R a_{\gamma-1,t-1+\gamma}, \quad \gamma \in \{\underline{\gamma}, \dots, \gamma_d - 1\}, \quad (11)$$

where $a_{\gamma,t+\gamma}$ denotes asset holdings by the end of period $t + \gamma$ at age γ — assuming no initial wealth at age $\underline{\gamma} - 1$: $a_{\underline{\gamma}-1,t-1+\underline{\gamma}} = 0$. Asset holdings are limited at each age by credit constraints

$$a_{\gamma,t+\gamma} \geq -\theta \frac{w_{\gamma+1,t+\gamma+1}}{R}, \quad \gamma \in \{\underline{\gamma}, \dots, \gamma_d - 1\}. \quad (12)$$

¹⁹We assume that agents die before their children enter into old age: $\gamma_d < \bar{\gamma} + \gamma_n$.

Fertility constraints. Fertility policies require that

$$n_t \leq n_{\max,t}, \quad (13)$$

$n_{\max,t}$ captures fertility policies at every date t . If at date t , agents can freely choose fertility, then $n_{\max,t} \rightarrow \infty$. In our experiments, fertility policy is unconstrained until date t_0 , and constrained thereafter by a sequence of $\{n_{\max,t}\}_{t \geq t_0}$.

Solution. Agents born at date t optimally choose a sequence of consumption $\{c_{\gamma,t+\gamma}\}_{\gamma \in \{\underline{\gamma}, \dots, \gamma_d\}}$, a level of fertility ($n_{t+\gamma_n}$) and human capital investment for their children ($h_{t+\gamma_n}$) in order to maximize their intertemporal utility $U(t)$ (Eq. 10), subject to a sequence of instantaneous budget constraints (Eq. 11), credit constraints (Eq. 12), and fertility constraints (Eq. 13). This characterizes consumption dynamics across age, as well as the dynamics of fertility and human capital $\{n_t, H_t\}_{t > 0}$ given initial conditions $\{n_0, H_0\}$. Details of the solution are provided in Appendix D.²⁰

4.2 Data and Calibration

Timing. Agents live for 20 periods, where a period lasts 4 years. They start working in the 6th period (ages 21-24) and have children in the 7th (ages 25-28)—in line with the data.²¹ They enter old age in period 16 corresponding to ages 61-64, the age at which males retire in China. Figure D.1 in Appendix D summarizes the timing and patterns of income flows and transfers, at each age of the agent’s life.

Endogenous variables prior to 1970 are assumed to be at a steady-state characterized by optimal fertility and human capital $\{n_{ss}; H_{ss}\}$. The calibrated parameters are summarized in Table 3 (details in Appendix D). Data used in the calibration are described in Appendix A.

Table 3: Calibration of Model Parameters

Parameter	Main Target (Data source)	Value
$R - 1$ (annual)	Average real interest rate, 1979-2013 (Bai et al. (2006)/CEI/NBS/PBOC)	5.3%
g_z (annual)	Real wage growth (UHS)	6.1%
α	Mankiw, Romer and Weil (1992)	0.37
v	Fertility in 1964-1969; $n_{ss} = 2.92/2$ (Census)	0.58
ω	Transfer to elderly w.r.t the number of siblings (CHARLS)	0.65
β (annual)	Age-saving profile in 1986 (UHS)	0.99
ψ	Age-saving profile in 1986 (UHS)	9%
θ	Age-saving profile in 1986 (UHS)	0%
ρ	Education expenditures across ages in 2002 (CHIP)	0.2
e_γ	Labour income by age in 1992 (UHS)	See Fig. 6 and 7
ϕ_γ	Compulsory education expenditures across ages in 2002 (CHIP)	and details in
$\phi_{\gamma,h}$	Discretionary education expenditures across ages in 2002 (CHIP)	Appendix D

Technology. The real growth rate of disposable income of Chinese urban households averages at a high rate of 7.3% over the period 1982-2014 (CEIC data). This rate of growth is an upper-bound


²⁰The model can be solved analytically if the credit constraints are not binding for ages $\gamma \geq \gamma_n$ (see Appendix D) — yielding a similar set of equations capturing the dynamics of fertility and human capital accumulation as in the model of Section 3; the model can otherwise be solved numerically.

²¹The average age of parents at first birth is 25.5 years in 1965-1970 and vary between 25 and 27 years until 1990 (Census).

for productivity growth g_z , as wage growth occurs partly endogenously through human capital accumulation. To estimate the rate of growth of g_z , we use individual income data from UHS over the period 1992-2009, estimating the average real wage growth over the period controlling for education (see Appendix D for details). On an annual basis, we obtain $g_z = 6.1\%$. The technological parameter α is set to 0.37 — in line with estimates of production functions in the empirical growth literature (Mankiw, Romer and Weil (1992) and Sianesi and van Reenen (2000)).²²

Age-Income Profile. We calibrate the experience parameters $\{e_\gamma\}_{\gamma \geq \gamma}$ to labour income by age group, provided by UHS data. The first available year for which individual labour income information is available is 1992. Calibrating the (pre-policy) initial income profile to 1992 data is sensible as human capital levels of the working-age population have not been affected by fertility controls (chosen by ‘non-treated’ parents). The age-income profile in 1992 is displayed in Figure 6.

Figure 6: Age income profiles in 1992 and 2009. Model vs. Data.



Notes: This figure plots the model-implied labour income profiles by age in 1992 and 2009 and its data counterpart. Data source: UHS, 1992 and 2009. Wages includes wages plus self-business incomes. The profile in 1992 is used to calibrate experience parameters $\{e_\gamma\}_{\gamma \geq \gamma}$. Parameter values for the model’s simulations are provided in Appendix D.

Real Interest Rate. In the spirit of Curtis et al. (2015) (see also Song et al. (2015)), we assume that the rate of interest R_t faced by households is defined by: $R_t = \lambda_t R_t^d + (1 - \lambda_t) R_t^K$, where R_t^d denotes the deposit rate which is controlled by the government and R_t^K denotes the return to capital implied by the marginal product of capital; λ_t measures the fraction of financial wealth of households in the form of deposits, which hovers between 70% and 90% in our data. Using data on R_t^d , R_t^K and λ_t , we compute the average real rate faced by households over the period 1979-2013. The resulting value of 5.3% is used to calibrate R (see Appendix D for details of the calibration of R).

Fertility, demographic structure and policy implementation. The targeted initial fertility rate n_{ss} is the one of urban households prior to 1970—when families were unconstrained. We use

²²Using Eq. 9, one can also compute α for a given ω by looking at the ratio of education expenditures per child of twins versus an only child (above 15). This method leads to an estimate of 0.39, which is very close to our calibrated value.

the average fertility over the period 1964-1969, equal to 2.92, to calibrate the initial steady-state and therefore select the preference parameter for children, v , to target $n_{ss} = n_{t < 1970} = \frac{2.92}{2}$. While the one-child policy became fully effective starting the 1980s, the policy also constrained households who started to conceive in the 1970s—accounting for the progressive decline in the 1970s as discussed in Section 2, and detailed in Appendix B. In our calibration, the one-child policy thus reduces fertility progressively during the 1970s, such that, taking cohorts to be born every year, fertility constraints ($n_{\max,t}$ for $1970 \leq t \leq 1980$) vary to match the fertility observed in the data over this period. For any date post-1980, fertility is constrained by the one-child policy: $n_{\max,t} = \frac{1}{2}$ for $t > 1980$.

We set the initial population distribution in 1964 to match the size of each age group above 17 years old in the Census 1982, age-bins (17-20, 21-24, ..., 77-80).²³ This makes sure that the composition effects driving aggregate saving are consistent with the population composition when the one-child policy is implemented. From this initial distribution, the population of each age group evolves in line with the path of fertility in the model and the data.²⁴

Old age support. Two parameters govern transfers to parents, ψ and ω . The first captures the degree of altruism towards parents in the economy; the latter captures the propensity to free-ride on the transfers provided by one’s siblings. We first estimate ω empirically.

Estimation of ω and validation of the transfer function. CHARLS provides data on transfers from a given child to his/her parents for the year 2008. Using variations in the amount of transfers to parents with different number of children, we estimate the log-transformation of the transfer function $\psi \frac{n^{\omega-1}}{\omega} w$. Details and results of the estimation are provided in Appendix D (Table D.2).

The amount of transfers (per offspring) given to parents is found to be decreasing with the number of siblings the offspring has, and increasing with the offspring’s income with an elasticity close to 1—validating empirically our transfer function. The elasticity ($\omega - 1$) of transfers to the number of children is estimated to -0.35. Thus, we set $\omega = 0.65$.

Measuring ψ . The parameter ψ measures the degree of altruism towards parents, linked to the overall level of transfers towards the elderly. Direct measurement of ψ based solely on measured transfers from CHARLS gives a very low value for ψ , around 4 – 5% for $\omega = 0.65$.²⁵ Such a low value does not square with the Census evidence where family support is reported to be the main source of income of elderly (Figure 2). Transfers measured in the data are likely to be underestimated. It does not include many forms of ‘non-pecuniary transfers’—in-kind benefits such as coresidence and health care—and CHARLS does not report most pecuniary transfers *within* a household in the case of coresidence. Section 2 documents how coresidence with children is a primary form of living arrangement for the elderly. Any transfer that provides insurance benefits to the elderly should in principle be taken into account—as they determine saving decisions for middle-aged adults. Importantly, if one takes pecuniary transfers towards parents living in another city from CHARLS (2011), one obtains a value of $\psi = 8\%$ — more in line with our calibrated value. These transfers are arguably a better proxy since in-kind benefits and mis-measured pecuniary transfers within households become less of an issue when parents live far away. Given the difficulty in accurately measuring ψ from the data, our preferred strategy discussed below is to calibrate it to match the age-saving profile in 1986.

²³The size of the age groups above 60 in 1964 (bins 61-64, ..., 77-80) remains undetermined, as above 80 in 1982. This is unimportant however for our purposes as these agents do not make human capital decisions for the later cohorts, and also because we focus on aggregate saving starting 1982, at which point they are no longer alive.

²⁴Our model fits the distribution of population in the later years reasonably well (see Appendix D). However, it predicts age-groups of older individuals larger than in the data as it does not feature mortality before age γ_d .

²⁵Wages of children, not observed in CHARLS (2008) can be imputed based on children’s characteristics. Transfers range from 4% (4 or more siblings) to 10% (only child) of the wages of individuals 42 – 54 years old, yielding a value of $\psi = 4 - 5\%$.

Parameters $\{\beta, \psi, \theta\}$ and education parameters $\{\rho; \phi_\gamma; \phi_{\gamma,h}\}_{\gamma \in \{\gamma_n, \dots, \gamma_n + \gamma_e\}}$. Our calibration strategy jointly determines the parameters $\{\beta, \psi, \theta\}$ and the education parameters $\{\rho; \phi_\gamma; \phi_{\gamma,h}\}_\gamma$ to best match the age-saving profile in 1986 (UHS data) while targeting education expenditures observed in 2002 (CHIP data) — 1986 (resp. 2002) is the first year for which we can measure saving by age (resp. education costs by age together with their decomposition between compulsory costs and discretionary costs).

Education expenditures observed in 2002 can be decomposed between compulsory costs (tied to parameters ϕ_γ) and discretionary costs (tied to parameters $\phi_{\gamma,h}$).²⁶ The fraction of wage income spent on compulsory education costs at a given age pins down the parameters $\{\phi_\gamma\}_{\gamma \in \{\gamma_n, \dots, \gamma_n + \gamma_e\}}$. As discretionary costs are very close to zero up to the age 10 of the child (Figure 7), we set $\phi_{\gamma,h} = 0$ for $\gamma \leq 8$ (age 29-32).²⁷ This ensures that, for the parameter values considered, education choices can be expressed analytically as the credit constraint is not binding when parents pay the discretionary costs (see Appendix D). Based on this analytical expression, we show that for each value of the parameter ρ , there is a unique combination of the parameters $\{\phi_{\gamma,h}\}_{\gamma \in \{\gamma_n, \dots, \gamma_n + \gamma_e\}}$ such that the rate of change of discretionary costs between two ages matches its data counterpart in 2002. For a given ρ , the parameters $\{\phi_{\gamma,h}\}_\gamma$ are thus set to match the shape of discretionary education costs by age — their overall level cannot be matched independently as it depends on the education choice of each generation of parents and on all the other parameters.

Having set the education costs parameters $\{\phi_\gamma; \phi_{\gamma,h}\}_\gamma$, we search for the remaining parameters $\{\beta, \psi, \theta, \rho\}$ over a grid $\mathbf{\Gamma}$ such that the model predicted age-saving profile in 1986 and the levels of discretionary education spending by age in 2002 are as close as possible from their data counterpart. More specifically, we search for parameters $\{\beta, \psi, \theta, \rho\} \in \mathbf{\Gamma}$ to minimize the following distance:

$$\min_{\{\beta, \psi, \theta, \rho\} \in \mathbf{\Gamma}} \left[\sum_{\gamma=\underline{\gamma}}^{\gamma_d} \lambda_\gamma^s \left| s_{\gamma,1986}^m(\beta, \psi, \theta, \rho) - s_{\gamma,1986}^d \right| + \sum_{\gamma=\gamma_n}^{\gamma_n + \gamma_e} \lambda_\gamma^{educ} \left| educ_{\gamma,2002}^m(\beta, \psi, \theta, \rho) - educ_{\gamma,2002}^d \right| \right]$$

where $s_{\gamma,1986}^m$ (resp. $s_{\gamma,1986}^d$) is the model predicted saving rate at age γ in 1986 (resp. the saving rate at age γ in the 1986 data); $educ_{\gamma,2002}^m$ (resp. $educ_{\gamma,2002}^d$) is the model predicted discretionary education spending as a share of wage at age γ in 2002 (resp. the discretionary education spending as a share of wage at age γ in the 2002 data); λ_γ^s and λ_γ^{educ} are weights on different age groups summing to one and reflecting their respective income share.


Intuitively, the parameter θ largely determines the saving rate at age 21-24—resulting in a very low value of θ . The value of the discount rate β mostly determines the aggregate saving rate, while ψ affects the overall shape of the profile — the amount of savings by individuals in their fifties and the corresponding dissavings in old age. Our combination of parameters gives a reasonable fit of the model-implied age-saving profile in 1986 with that of the data (Figure 8, upper panel).²⁸ The last parameter ρ guarantees that the level of education spending stays in line with the data given all the other parameters — the whole combination of education parameters $\{\rho; \phi_\gamma; \phi_{\gamma,h}\}_\gamma$ fitting data

²⁶These estimates based on education expenditures represent a lower bound for the cost of children, as other forms of transfers (food, co-residence,...) are largely omitted. But, unlike education costs, these expenditures are difficult to break down into amounts solely related to children.

²⁷Education costs are paid until age 14 (age 53 to 56 years)— $\gamma_e = 7$.


²⁸As our sensitivity analysis shows (see Appendix D), taking $\psi = 4\%$ from direct estimates (CHARLS) significantly distorts the profile. Lower transfers to the elderly increases significantly the saving of the middle-aged — as lower receipts of transfers from children bid the middle-aged to save more. This larger wealth accumulation also leads to larger dissaving of the old compared to the data.

Figure 7: Education expenditures per child by age of parents in 2002. Model vs. Data.



Notes: This figure plots the model-implied education profiles by age of parents in 2002 and its data counterpart (in % of income). The left-panel shows compulsory education costs per child and the right panel shows discretionary education costs. Parameter values for the model's simulations are provided in Table 3 and detailed in Appendix D. The data counterpart is computed using CHIP 2002 (see Appendix A for details on education expenditures data).

Figure 8: Age-saving profile in 1986 and 2009. Model vs. Data.



Notes: This figure plots the model-implied age-saving profile in 1986 and 2009 and its data counterpart. Parameter values for the model's simulations are provided in Table 3 and detailed in Appendix D. The data counterpart is estimated using UHS data (see Appendix E.1 for details on the estimation procedure).

on education spending in 2002 extremely well (Figure 7). The minimization leads to the following parameter values: $\beta = 0.99$ (annual basis); $\psi = 9\%$; $\theta = 0\%$; $\rho = 0.2$ — the corresponding education costs $\{\phi_\gamma; \phi_{\gamma,h}\}_\gamma$ parameters being shown in Appendix D. The discount rate β is admittedly high though still in the ballpark of related papers.²⁹ As household saving is fairly high, it remains difficult to match its level without a high discount rate in a model without uncertainty and precautionary saving. Credit constraints are found to be very tight, in line with the low dissavings of young households and the low level of household debt in China.³⁰ Most importantly, the resulting value for the transfer parameter ψ is in line with Banerjee et al. (2014) and in line with data on pecuniary transfers towards parents living in another city. Sensitivity analysis with respect to the main parameters of the model is relegated to Appendix D.

4.3 Results

We now investigate the impact of fertility policies in our quantitative model on various outcomes, from aggregate implications to micro-level predictions.

4.3.1 Aggregate saving and human capital accumulation

Aggregate saving. Figure 9 displays the aggregate household saving rate in the years following the fertility policies in the model and in the data. In our baseline simulation, the aggregate saving rate increases by 11.6% over the period 1982-2014, about 60% of the increase in the data. This is an *upper-bound* of what can be attributed to the policy change—as the natural fertility rate might have fallen since 1982 and thus raised saving independently of the policy. Section 5.2 discusses counterfactual fertility and saving in the absence of the policy. Our model also predicts a fall in aggregate saving in the coming years as a result of compositional shifts (macro-channel), whereby the only child generation ages and old dissavers account for a larger share of the population. In our simulation, we decompose the effect on saving driven by the ‘micro-economic channel’ (transfer and expenditure effects) and by the ‘macro-economic channel’ (composition effects). To do so, we simulate the increase in aggregate saving due to changes in the saving rate across ages while keeping the population composition fixed to its 1982 counterpart. This isolates the effect due to the ‘micro-economic channel’ (dotted line on Figure 9)—the remaining increase in aggregate saving being due to composition effects. Our decomposition shows that the ‘micro-economic channel’ is quantitatively large, contributing to more than 60% of the 11.6% increase in the saving rate predicted by our model.


It is reassuring that the dynamic of the saving rate is not very sensitive to different values of ψ — a 11.6% rise over the period 1982-2014 in the baseline calibration ($\psi = 9\%$) compared to a 10% rise in the case of low transfers ($\psi = 4\%$). The predicted change in the aggregate saving rate is of similar order of magnitude because the two main channels governing aggregate saving turn out to be more or less offsetting when varying ψ : a higher ψ makes the ‘micro-channel’ stronger owing to a greater importance of transfers; however, the ‘macro-channel’ is dampened since composition effects on saving are weaker when differences in saving rates among age groups are less pronounced. The predicted rise in aggregate saving is thus comparable despite different age-saving profiles across calibrations.³¹

²⁹Song, Storesletten, and Zilibotti (2011), Banerjee et al. (2014) and Curtis et al. (2015) use values between 0.99 and 1.

³⁰The lack of consumer credit and mortgage markets, and the very low levels of household debt in China (about 10% of GDP in 2008) warrants a choice of a low θ to strongly limit the ability of young households to borrow against future income. Our baseline calibration gives $\theta = 0$ since the saving rate of the 21-24 age group is slightly positive in 1986. It is slightly negative in later years but results are not sensitive to θ as long as it is not too large. See Appendix D for sensitivity analysis.

³¹See Appendix D for sensitivity analysis with respect to ψ . Note that in order to match the level of aggregate saving

Figure 9: Aggregate Household Saving Rate: Model vs. Data



Notes: Data source: CEIC Data (using Urban household Survey, UHS). The model implied aggregate saving rate simulates the fertility policies using the calibration of Table 3.

Human capital accumulation. Due to the quantity-quality trade-off, our model predicts an increase in the level of human capital in the economy following the policy. Quantitatively, the level of human capital for an only child is 53% higher than the level of his/her parents (with two siblings)—translating into a wage increase of about 17% for the only-child generation compared to their parents. This has to be weighed up against a 49% increase in the number of years of schooling of the only-child generation compared to their parents—even though comparison between model and data is not straightforward.³² As a consequence, the model generates *endogenously* a portion of the flattening of the age income profile observed in the data in 2009. The increase in human capital of the only-child generation explains a significant fraction of the faster wage increase of young adults (Figure 6).

Using cross-sectional comparison between twins and only child born in the 1980s, the model predicts that a twin reaches a level of human capital 24% lower than an only child. Note that the human capital difference between an only child and a twin is comparable to the overall effect of the policy if the natural fertility rate is around 2. Differences in education spending and attainment between twins and only child are testable implications that motivate our subsequent empirical strategy.

The policy response of human capital is also critical for assessing the quantitative impact of a fertility change on aggregate saving. As shown in Section 3.3, the degree of substitution between quantity and quality determines the extent of the fall in expected transfers—and thus the strength of the saving-response through the transfer channel.³³

with a lower ψ , one needs to reduce also the discount rate β . With a $\beta = 0.98$ — all other parameters being identical, the increase in aggregate saving over the period 1982-2014 is 10%.

³²According to Barro and Lee (2010), for China as a whole, the generation of only child aged 25-34 in 2010 has on average 8.7 years of schooling compared to 5.8 for their parents. The number of children achieving secondary education being multiplied by almost 3 over 1980-2010. See Li et al. (2013) for similar numbers on urban households only.

³³Alternative calibrations (of decreasing returns to transfers and/or education for instance) can generate a much stronger human capital response to the policy and thus a very low saving response—micro-channels being limited to the ‘mouth to feed’ expenditure channel. Note also that human capital accumulation also shifts the distribution of income across age groups, and in turn impacts aggregate saving (income composition channel). This channel rises in magnitude post-2010 when the generation of only children exerts a greater impact in the economy in terms of their higher income and saving.

4.3.2 Micro-implications: age-saving profiles

Age-saving profiles: model versus data. The bottom panel of Figure 8 shows age-saving profiles in 2009 in the data and in our model. This profile has to be compared with the one in the earlier period of our sample in 1986 that is used in the calibration (upper panel). The data reveals some evolution in the age-saving profile between 1986 to 2009. There has been an upward *shift* in the age-saving profile for all age groups but the youngest ones, as well as a change in the *shape* of the profile. The latter is characterized by a flattening of the saving profile for the middle aged (30-50), which contrasts with the more hump-shaped profile in 1986.³⁴ One can mark that the model can generate the upward *shift* of the profile over the period except for the oldest above 60. The rise in saving rates results from both a fall in expenditures on children and a fall in expected future receipts of transfers. The model’s predicted profile for 2009 is quantitatively in line with the data with the exception of the oldest age-groups. As these people were unaffected by the policy, our model alone cannot explain the rise in their savings. Abstracting from this group, the model also captures part of the shape of the profile in 2009. In particular, it predicts a fairly high saving rate for individuals in their early 30s even though not as high as in the data: in 2009, they are the most impacted by the policy—because they are only child themselves and therefore take on the brunt of the burden of supporting their parents later, but also because they are subject to the one-child policy and therefore expect to receive less transfers from their own only child.


Age-saving profiles: twins versus only child. As shown in our theoretical analysis, one can relate the time-series change in saving rates to the cross-sectional difference in saving rates between a twin-household and an only-child household. Figure 10 (left panel) plots the difference in saving rates at a given age between parents of an only child and parents of twins as predicted by the model for a 2006 cross-section of individuals: $(s_{\gamma,2006} - s_{\gamma,2006}^{twin})_{\gamma=\{\underline{\gamma},\dots,\bar{\gamma}\}}$. As stated by Lemma 2, this difference is a lower bound of the overall time-series change in saving rates implied by the policy—if the natural fertility rate stays above 2. For comparison purposes, the right panel of Figure 10 shows the overall change in saving rates across age groups from the initial steady-state in 1970 to 2006.³⁵ As expected, the predicted increase in saving rate over time is quantitatively larger than the difference of saving rates between parents of an only child and parents of twins. However, the very same patterns emerge: only child households save more across all age groups, even after children have departed from the household—when the expenditure channel is no longer in operation. The difference of saving rates between parents of an only child and parents of twins is at the heart of our identification strategy developed in Section 5 — investigating this difference in the data provides a clear test of the quantitative properties of our model.

Age-saving profiles: comparison with a standard OLG Model. Figure 10 displays also the difference in saving rates across ages in a standard OLG model without old-age support — differences in saving rates between parents of an only child and parents of twins (left panel) and differences over time over the period 1970-2006 (right panel). In this standard OLG model, only the *expenditure*

³⁴Chamon and Prasad (2010) and Song and Yang (2010) show that saving rates augmented more for young households over the period. While our profiles show a significant increase for households in their thirties, an important difference between our saving profiles as estimated from the data and those of Chamon and Prasad (2010) and Song and Yang (2010) is that young (childless) adults did not see a rise in saving rates. The difference comes from our correction for the biases associated with multi-generational households (see Coeurdacier, Guibaud and Jin (2015)).

³⁵We use the prediction in 2006 as the data counterpart in our sample of twins cover the years 2002-2009. Results using 2009 as final year are very similar.

Figure 10: Difference in saving rates by age between parents of an only child and parents of twins (left-panel) and between 2006 and 1970 (right-panel). Model Predictions.



Notes: The left panel of this figure plots the model-implied difference in saving rates between parents of an only child and parents of twins in 2006 at different ages: $(s_{\gamma,2006} - s_{\gamma,2006}^{twin})_{\gamma=\{\underline{\gamma},\dots,\bar{\gamma}\}}$. The right-panel plots the model-implied change in saving rates across age-groups over the period 1970-2006: $(s_{\gamma,2006} - s_{\gamma,1970})_{\gamma=\{\underline{\gamma},\dots,\bar{\gamma}\}}$. Two cases considered: our baseline calibration and standard OLG model in which old age support and human capital accumulation are absent. Parameter values provided in Table 3.

channel is operative.³⁶ In the absence of old-age support, the standard OLG model predicts a much smaller difference in saving rates across all ages. Thus, the transfer channel appears to be necessary for matching quantitatively the change in saving rates over time and the cross-sectional difference in saving between parents of an only child and parents of twins. Another important discrepancy between the two models concerns individuals in their late 50s (resp. early 30s). Due to consumption smoothing over the life cycle, lower expenditures on children release more resources for consumption when children are no longer living in the household (resp. when education costs are very low as in the children’s early ages). Thus, the standard OLG model predicts lower saving rates for these age groups in households with fewer children, while our model predicts higher ones due to the *transfer channel*. In particular, this channel can be identified by investigating the saving behavior of parents in their 50s once children have left the household—as done in Banerjee et al. (2010, 2014). Importantly, the magnitude of the transfer channel predicted by our model is very close to their empirical evidence. Using the partial implementation of fertility restrictions in the 1970s as an identification strategy, their estimation compares the saving behavior of individuals in their early 50s to individuals in their early 60s in 2008: in line with our quantitative estimates, they find that the latter save on average about 10% less than the former.³⁷

³⁶Education costs per child ϕ_{γ} are kept constant but human capital is fixed and transfers to elderly are set to zero. Similar patterns emerge if old-age support is independent of the number of children.

³⁷One can mark that our benchmark model falls short of explaining the sharp increase in saving rates of older workers and retirees over this period. Banerjee et al. (2010, 2014) suggests that our model predicts the appropriate variations between treated and non-treated households. Yet, saving rates increased for both groups beyond what can be explained by fertility restrictions.

5 ‘Twin’ Tests and Counterfactuals: Model vs. Data

Section 3 showed how one can identify theoretically the micro-channel by comparing two-children (twin) households to only-child households. Using this theoretical analysis as guidance, we estimate a ‘twin effect’ from the data and, using the ‘twin’ experiment in the quantitative model, we compare various outcomes between model and data. Our strategy is to compare the decisions of parents of an only child to decisions of parents with an exogenous extra-child (twins) under the one child policy. The mere presence of the one-child policy allows us to circumvent some identification issues when using the birth of twins as an exogenous fertility shock. For instance, without the policy, twinning is more likely to occur when families have more kids and this preference for fertility could be correlated with parental decisions. Under the one-child policy (post-1980), identification becomes cleaner as households have either one child or randomly two (twins).³⁸ One may still question the validity of using twins as exogenous deviation of fertility—in the event that twinning is not random, for instance fostered by ‘artificial’ fertility methods whose adoption may be correlated with the propensity to save. We endeavor to address this concern. The important thing to note is that identification based on twins born under the one-child policy is of independent value—particularly for providing an out-of-sample check to our model predictions—and this is precisely how it should be viewed.

5.1 Estimates of the ‘Twin Effect’

A detailed description of the data used is provided in Appendix A. A data limitation is that one observes children (twins or only child) only when (1) residing in a household, (2) when residing outside but remaining financially dependent, or (3) in the years following their departure from the household using the panel dimension of UHS.³⁹ Ideally, one would like to observe the saving behavior of parents in their fifties, after the children have departed. This limitation means that the ‘transfer channel’ can only be inferred from the fewer observations of older parents still living with their children, or from parents whose children had just left the household—rather than using the whole set of observations of parents in their fifties living alone.

Household saving. The first set of regressions estimates the impact of twins on household saving rate. It uses the whole sample in UHS (1986 and 1992-2009), which includes households that had children both before and after the implementation of the one-child policy. We consider only households with resident children below the age of 18 (or 21 as a robustness check), as otherwise consumption, income and saving of the household include those of the potentially employed children. The following regression is performed for a household h living in province p at a date $t = \{1986, 1992, \dots, 2009\}$:

$$s_{h,p,t} = \alpha_t + \alpha_p + \beta_1 D_{h,t}^{\text{Twins born} > 1980} + \beta_2 D_{h,t}^{\text{Twins born} \leq 1980} + \gamma Z_{h,t} + \varepsilon_{p,h,t}, \quad (\text{R1})$$

where $s_{h,p,t}$ denotes the household saving rate of household h (defined as the household disposable income less expenditures over disposable income); α_t and α_p are respectively time and province fixed-

³⁸While the policy was effective starting 1980, it has also affected households who started to procreate in the 1970s as it takes time to conceive children (see discussion in Section 2). These households having their first child in the 1970s might end up with more than one child but can be nevertheless constrained when the policy is implemented. Thus, an identification based on before/after the shock comparison is likely to fail. Our identification strategy relying on comparing the behavior of twin parents versus parents of only child under the policy regime (post-1980) also circumvents this difficulty.

³⁹Family composition and the number of children are in general unobserved in UHS when children are financially independent and live outside of the household. The panel dimension (households observed for 3 consecutive years) provides some observations of households where children have departed.

effects, $D_{h,t}^{\text{Twins born} > 1980}$ is a dummy that equals 1 if the twins are born after the full implementation of the one-child policy (post 1980), $D_{h,t}^{\text{Twins born} \leq 1980}$ is a dummy that equals one if twins born before 1980 are observed in a household, $Z_{h,t}$ is a set of household level control variables and $\varepsilon_{p,h,t}$ is the residual. β_1 measures the effect of having twins under the one-child policy regime (post-1980) and is the coefficient of interest: it measures the effect on the household saving rate of having twins instead of an only child. β_2 is less relevant for our purpose — it measures the effect on the household saving rate of giving birth to twins before 1980 and is more difficult to interpret since the one-child-policy was not binding and there might be some selection into twinning.

Table 4: Household Saving Rate: Twin Identification

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	Sav. rate	Sav. rate	Sav. rate	Sav. rate	Sav. rate inc. educ. transfers	
Oldest child	Up to 18y	Up to 18y	Up to 21y	Up to 18y	Up to 18y	Up to 21y
Sample UHS	1986 and 92-09	1986 and 92-09	1986 and 92-09	2002-2009	2002-2009	2002-2009
Type of household	All	All	All	Nuclear only	Nuclear only	Nuclear only
Twins born > 1980	-0.0572*** (0.0106)	-0.0540*** (0.00971)	-0.0566*** (0.00937)	-0.0717*** (0.0123)	-0.0839*** (0.0127)	-0.0789*** (0.0122)
Twins born ≤ 1980	0.0177 (0.0173)	0.0143 (0.0162)	0.0197 (0.0142)			
Benchmark Controls	YES	YES	YES	YES	YES	YES
Additional Control (1)	NO	YES	YES	YES	YES	YES
Additional Control (2)	YES	YES	YES	NO	NO	NO
Observations	84,403	84,403	100,236	41,746	41,706	50,439
R-squared	0.074	0.173	0.165	0.184	0.184	0.185
Years Dummies	YES	YES	YES	YES	YES	YES
Province Dummies	NO	YES	YES	NO	NO	NO
Prefecture Dummies	NO	NO	NO	YES	YES	YES

Notes: Data source: UHS (1986, 1992-2009). We take one observation per household. Outliers with saving rate over (below) 85% (-85%) of income are excluded. Controls include average age of parents, mother’s age at first birth, and child’s age. Additional Control (1) includes household income in addition to the benchmark controls, and Additional Control (2) includes a dummy for the multigenerational structure of the family. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Columns (5) and (6) include education transfers to children living in another city as part of consumption expenditures when computing household saving.

Columns 1-3 in Table 4 display the coefficient estimates of the impact of twins on household saving rate before and after the policy implementation. The estimated coefficients on $D_{h,t}^{\text{Twin born} > 1980}$ show that under the one-child policy, households with twins saved (as a share of disposable income) on average 5 to 6 percentage points less than household with an only child. The magnitude is similar under different specifications and across samples.⁴⁰

Columns 4-6 report regression results for a restricted sample of nuclear households (unigenerational). These households had only one incidence of births—either bearing an only child or twins. The advantage of pooling all households that are unigenerational is that the same demographic composition (up to the presence of twins) applies to all households —making this exercise the closest to our theoretical framework. Unlike the full sample in regression (R1), all households are having chil-

⁴⁰In Column 1, household income is excluded because it could be an outcome variable—household members with a large number of children may decide to work more to meet higher expenditures, or, decide to reduce the labor supply of mothers. Column 2 controls for household income. Column 3 includes all children up to the age of 21 years old.

dren after the implementation of the one-child-policy.⁴¹ Households with twins have on average a 7 percentage-points lower saving rate than those with an only child (Column 4). The effect estimated in the cross-section of (fully) treated unigenerational households gives results fairly close to the estimates using the whole sample of households (Columns 1-3). In Columns 5-6, we compute an alternative and more accurate measure of the saving rate by incorporating education transfers to children residing outside of the household as part of household expenditures (only available in the sample starting in 2002). The more precise measure of saving rate gives a larger twin effect: households with twins save about 8 percentage-points less than those with an only child. In a nutshell, our results show that having (exogenously) one more child under the one-child policy reduces saving rates by at least 5 percentage-points and up to 8 percentage-points.

Identifying the transfer channel. One may argue that the results on saving are driven entirely by the extra costs of having twins compared to an only-child, as one cannot disentangle the ‘expenditure channel’ from the ‘transfer channel’ in the previous regressions. We use two different strategies to provide evidence for the relevance of the ‘transfer channel’ — one based on parental age, and one that identifies a specific ‘twin effect’ on saving after their departure from the household.

The ‘transfer channel’ becomes more relevant in one’s older age as shown in Section 4.3.2. At the same time, it should primarily affect non-education related expenditures. We test whether there is a differential twin effect for older parents (above 45), and particularly so for expenditures excluding education. Results are shown in Table 5 using the sample of nuclear households (unigenerational). The first observation is that savings of twin-households compared to that of only-child households are smaller — but even more so for parents above 45 (Columns 1-2). Furthermore, expenditures excluding education are higher for twin households and again particularly so for older parents (Columns 3-4). This is very suggestive that the ‘transfer channel’ is in operation.

Table 5: Savings and expenditures for different age groups: Twin identification

VARIABLES (in % of household income)	(1) saving rate	(2) saving rate	(3) Non-education exp.	(4) Non-education exp.
Twins	-0.0839*** (0.0127)	-0.0655*** (0.0137)	0.0360*** (0.0132)	0.0195 (0.0144)
Twins with parents ≥ 45		-0.110*** (0.0347)		0.0841** (0.0338)
Controls	YES	YES	YES	YES
Observations	41,706	41,706	25,716	25,716
R-squared	0.184	0.185	0.170	0.170
Years Dummies	YES	YES	YES	YES
Prefecture Dummies	YES	YES	YES	YES

Notes: Data source: UHS (2002-2009) for columns 1-2 and UHS (2002-2006) for columns 3-4 (decomposition of expenditures across different sectors including education is only available for the years 2002-2006). For columns 1 and 2, education expenditures include education transfers to children living in another city. Restricted sample of nuclear households are those with either an only child or twins up to the age of 18 years old. Outliers with saving rate over (below) 85% (-85%) of income are excluded. In columns 3-4 outliers with non-education expenditures above 150 % of income are also excluded. Controls include average age of parents, mother’s age at first birth, child’s age, and household income. In columns (2) and (4) dummy for parents above the age of 45. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

⁴¹The regression performed is for a household h living in prefecture p at date $t = \{2002, \dots, 2009\}$: $s_{h,p,t} = \alpha_t + \alpha_p + \beta D_{h,t}^{\text{Twins}} + \gamma Z_{h,t} + \varepsilon_{p,h,t}$.

Table 6: Saving differences between twins and only child: identification on ‘movers’

VARIABLES	(1)	(2)
	Sav. rate	Sav. rate
Oldest child	Up to 30y birth \geq 1980	Up to 30y birth \geq 1980
Adult twins left the household	-0.0920 (0.0728)	-0.0910 (0.0728)
Adult singleton left the household	0.0698*** (0.0117)	0.0708*** (0.0119)
Twins	-0.0498*** (0.00976)	-0.0546*** (0.0109)
Twins 18 to 30y		0.0189 (0.0236)
Singleton 18 to 30y		0.00127 (0.00284)
Observations	82,922	82,922
R-squared	0.171	0.171
Additional controls	YES	YES
Years Dummies	YES	YES
Province Dummies	YES	YES

Notes: Data source: UHS (1992-2009). We take one observation per household. Outliers with saving rate over (below) 85% (-85%) of income are excluded. The sample is restricted to households with either a singleton or twins in at least one of the survey waves. Controls include, in logs, the average age of parents, mother’s age at first birth, average child’s age and household income. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

To identify the ‘transfer channel’ as the main source of variation of saving rates across households with a different number of children, one would prefer to observe saving after the children have departed from the household and have become financially independent.⁴² The panel dimension of UHS partially allows for this, identifying a specific effect on parental saving on ‘movers’—households for which twins (or singleton) have left the household in between two surveys. Unfortunately, this is at the expense of the number of observations for identification as ‘movers’ constitute a small fraction of our sample of twins (about 20 observations).⁴³ Results are shown in Table 6 using the sample of households with children. Columns 1-2 show how savings of parents of twins and only child are affected once one (or two) child has left the household (the reference group being households with an only child residing in the household). For households with an only child, the saving rate is higher once the child has left—whereas it falls, if anything, for twins (although the coefficient is not statistically different from zero). Most importantly, households with an only child still save more than twin households once a child has left. Column 2 checks that our findings are not driven by the older age of ‘movers’.

‘Artificial’ twins. There is a concern that twins born after the one-child policy could potentially be ‘artificial’ or ‘man-made’ (Huang et al. (2016)). If true, this becomes an issue when families with ‘artificial’ twins have a different propensity to save/educate—after controlling for observable factors such as differences in household income, education, parents’ age, etc. In our urban sample, we do not observe significant deviations of twin births from the biological rate, neither before nor after 1980. This is consistent with Huang et al. (2016), who also do not find significant manipulation of twins for urban households.⁴⁴ We also conducted a series of robustness checks on income and saving differences


⁴²The ‘expenditure channel’, if anything, would tend to *raise* the saving rates of families with more children, once they have left—owing to consumption smoothing (see Figure 10).

⁴³Due to the lack of ‘movers’ in the twins sample, we have to consider households in which one or two children have left, rather than those in which both have left.

⁴⁴Huang et al. (2016) do not find evidence of twin manipulation in their urban sample of over a million observations, their

between only-child and twin households over time. If ‘artificial’ twin households were partly driving our empirical results, the difference between the two types of households would increase over time, as ‘artificial’ twinning technologies improve and become more accessible. Our investigation does not support this hypothesis.⁴⁵

Figure 11: Education Expenditures per child: Only Child vs. Twins



Notes: UHS (2002-2006), restricted sample of nuclear households. This figure displays the average education expenditure per child (as a share of total household income) by age of the child, over the period 2002-2006.

Quantity-Quality Trade-Off. A quantity-quality trade-off is immediately visible from the evidence in Figure 11: the per-capita education expenditure on a twin is significantly lower than on an only child—for children above the age of 15. The difference reaches almost 40% at age 20. One can confirm this finding by running the regression

$$\frac{\exp_{h,p,t}^{Educ.}}{n_{h,t}} = \alpha_t + \alpha_p + \beta D_{h,t}^{Twin} + \gamma Z_{h,t} + \varepsilon_{p,h,t}, \quad (R2)$$

for a household h at date $t = \{2002, \dots, 2006\}$, where $\frac{\exp_{h,p,t}^{Educ.}}{n_{h,t}}$ denotes the education expenditure household h spends on each child (as a share of household income) at date $t = \{2002, \dots, 2006\}$.⁴⁶

Results of regression (R2) are shown in Columns 2 and 4 of Table 7. For the sake of comparison, the impact of twins on overall education expenditures of the household is also shown (Columns 1 and 3). We find that education investment (*per* child) in twins is significantly lower than in an only child: while having twins significantly raise total education expenditures (as a share of household income) (Column 1), it reduces education expenditures spent on each child—by an average of 2.1 percentage points (Column 2). As conjectured, this trade-off mostly applies to older children (above 15), whose education attainment becomes more discretionary (Column 4).

findings being driven by the rural sample.

⁴⁵While there is a clear discontinuity between twin and non-twin household’s saving behavior around 1980 (echoing our regression results), the difference between their saving rates has not risen over time since then. Also, no such discontinuity occurred for the average household income level—which has been similar between twin and non-twin households (by first child birth) since 1970—nor for the number of observations of twin vs. non-twin households since 1970: the proportion over this period has stayed roughly constant.

⁴⁶Education expenditures are only available for the years 2002-2006 in UHS.

Table 7: Education Expenditures per Child: Twin identification.

VARIABLES (in % of household income)	(1)	(2)	(3)	(4)
	Education exp. total	Education exp. per child	Education exp. total	Education exp. per child
Twins	0.0648*** (0.0108)	-0.0215*** (0.00539)	0.0533*** (0.0101)	-0.00917* (0.00510)
Twins \geq 15			0.0277 (0.0225)	-0.0248** (0.0113)
Controls	YES	YES	YES	YES
Observations	31,513	31,513	31,513	31,513
R-squared	0.127	0.126	0.141	0.140
Years Dummies	YES	YES	YES	YES
Prefecture Dummies	YES	YES	YES	YES

Notes: UHS (2002-2006), restricted sample of nuclear households are those with either an only child or twins up to 21 years of age. Education expenditures include education transfers to children living in another city. Other controls include average age of parents, mother’s age at first birth, child’s age and household income. Outliers with saving rates over (below) 85% (-85%) of income are excluded. Robust standard errors in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 8: Education Attainment: Twin Identification (LOGIT)

VARIABLE (logistic regression)	Higher education		Academic high school		Technical high school	
	(1)	(2)	(3)	(4)	(5)	(6)
	estimate	odds ratio	estimate	odds ratio	estimate	odds ratio
Twins	-0.489*** (0.158)	0.613*** (0.0968)	-0.455*** (0.138)	0.635*** (0.0875)	0.269* (0.157)	1.308* (0.205)
Controls	YES	YES	YES	YES	YES	YES
Observations	15,313	15,313	15,313	15,313	15,313	15,313
Years dummies	YES	YES	YES	YES	YES	YES
Province dummies	YES	YES	YES	YES	YES	YES

Notes: UHS (2002-2009), restricted sample of nuclear households are those with either an only child or twins of ages 18-22 years old. Controls include child’s age, average age of parents, mother’s age at first birth, average parents’ education level, and household income. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The quantity-quality trade-off is also visible when looking at differences in education attainment. LOGIT regression results on dummies that measure the level of school enrollment (academic high school, technical high school and higher education) are displayed in Table 8. Comparing education attainment of twins versus only children (of age 18-22) over the period 2002-2009 indicates that twins are on average 40% less likely to pursue higher education than their only-child peers (Column 2), a quantitatively large effect. The reason is that twins are about 40% less likely to pursue an academic secondary education that prepares to university (Columns 4) and instead 30% more likely to attend a technical high school (Column 6).⁴⁷

Predictions of the ‘Twin Effect’: Model vs. Data. We turn to the simulated results of a twin experiment as predicted by our model (and discussed in Section 4.3), and juxtapose these results with

⁴⁷It is possible that twins are of lower quality compared to singletons—for example, by having lower weights at birth (net of family-size effects)—and parents may in turn invest less in their education and substitute investment towards their singleton offspring. The problem is less serious, however, when households are allowed only one birth as in China. Oliveira (2012) finds no systematic differences between singletons and twins.

Table 9: Twin Experiment: Model and Data

		Model			Data ^a
Saving rate		Only child	Twins	Difference	Difference
$\gamma = 7 - 8$	(25–32)	7.1%	5.8%	1.3%	-1.1 – 2.9% (<i>n.s.</i>)
$\gamma = 9 - 10$	(33–40)	21.3%	16.3%	5.0%	4.9 – 5.4% (***)
$\gamma = 11 - 12$	(41–48)	33.4%	25.4%	8.0%	7.1 – 10.4% (***)
$\gamma = 13 - 14$	(49–56)	40.1%	35.7%	4.4%	11.3 – 16.2% (**)
Education expenditures					
(% of wage income)					
$\gamma = 7 - 8$	(25–32)	2.5%	5.0%	-2.5%	-0.9% (<i>n.s.</i>)
$\gamma = 9 - 10$	(32–40)	6.0%	11.6%	-5.6%	-4.2% (***)
$\gamma = 11 - 12$	(41–48)	9.7%	17.3%	-7.6%	-9.8% (***)
Human capital		Only child	Twins	% Difference	
$(H_{1986} - H_{ss}) / H_{ss}$		53%	16%	$\left(\frac{H_{only} - H_{twin}}{H_{only}} \right) = 24\%$	

^aEstimates of the impact of twins on household saving rates and education expenditures for parents in the different 8 years age brackets are available on request. We control for five parents age brackets between 25 and 64 years old, and report the highest and lowest point estimates of the interaction between the ‘Twins born after 1980’ dummy and the five age brackets. The specifications for the saving rate regressions are similar to the ones in columns 1 to 4 of Table 5, and the specification for the education expenditures regression is similar to column 1 of Table 8. For the age bracket [49–56], we use the saving estimation based on ‘movers’ (columns 1 to 3 of Table 7). *** (resp. ** or *) for estimates different from zero at the 1% level (resp. 5% or 10% level). (*n.s.*) for estimates non-significant at 10%.

Notes: This table compares the saving rate, expenditures devoted to children and children’s human capital attainment for households with twins and those with an only child, under the baseline calibration in 2006, and in the data (where relevant).

empirical estimates. Table 9 reports model outcomes in 2006 for an individual with twins and an individual with an only child at various parental ages. The model predicts fairly close estimates on the differences between these individuals compared to data estimates until age 48.⁴⁸ The predicted saving rate at $\gamma = 9 - 10$ and $\gamma = 11 - 12$ are respectively 5% (4.9 – 5.4% in the data) and 8.0% (7.1 – 10.4% in the data) lower in households with twins than in households with an only child. Above age 48, once children have left, estimates from the data based on movers are less in line with our predictions, but arguably less precisely estimated (a 4.4% difference in the model against more than 10% in the data, even though for the latter, standard errors are large). When examining education expenditure differences (as a share of wage income), we observe that households with twins have 5.6% (4.2% in the data) higher total expenditures for $\gamma = 9 - 10$ and 7.6% (9.8% in the data) higher expenditures at $\gamma = 11 - 12$.⁴⁹ Our calibrated model suggests a 24% difference in human capital attainment between a twin and an only child— compared against a 40% smaller chance of accessing higher education in the data. The proximity of model and data estimates are reassuring since the model is not calibrated on twin household variables.

5.2 Counterfactuals

Empirical Counterfactual of a ‘two-children’ policy. Using the empirical estimates of the twin-effect on saving and human capital, one can back out the counterfactual aggregate saving rate if parents were having two children instead of one (‘two-children’ policy). The procedure to compute

⁴⁸In the data, we estimate the difference across bins of 8 years (comprising two age groups of 4 years) to preserve a sufficient number of observations for twins.

⁴⁹Due to the quantity-quality trade-off, parents of twins thus spend about 1.1 percentage points less on discretionary education per child (as a % of wages) at age $\gamma = 11 - 12$ — smaller than the 2% estimate in the data (Table 7).

the counterfactual involves estimating the age-saving profile and aggregate saving rate that would have prevailed in 2009 if all households were assumed to have two children after 1980, and to behave like parents with twins (using regression results based on twins).⁵⁰ Details of the methodology are provided in Appendix E.2.

Results are displayed in Table 10, which shows the decomposition of the overall effect of the policy on aggregate saving into contributions from the various channels. The counterfactual exercise indicates that the aggregate saving rate would have been between 6.0% and 6.8% lower if China had implemented a ‘two-children’ policy. These empirical estimates attribute to the one-child policy about a third of the increase in the aggregate saving rate since 1982. Importantly, the micro-channels explain most of the overall effect, and are significantly more important than the macro-channel conventionally emphasized. Note that such a strategy provides a lower-bound of the overall impact of the one-child policy—assuming that the natural rate of fertility would have stayed above 2 until recently.

Table 10: Empirical counterfactuals using estimates from twins regressions: aggregate effect under a two children policy.

	Aggregate saving rate	Additional effect
Aggregate saving rate 2009 (Census corrected)	29.89%	
Macro channels		
Age composition	28.27%	-1.62%
Education and income composition (22 to 29y)	28.64%	0.38%
Micro channels		
Education (child below/over 15y)	26.12%	-2.52%
Non-education (parents below/above 45)	24.69%	-1.43%
Additional transfer channel ^a		
(More conservative) 3% scenario	23.90%	-0.79%
(Less conservative) 6% scenario	23.12%	-1.57%
Total effect (3% scenario)		-6.0%
Total effect (6% scenario)		-6.8%
Total effect (model counterfactual in 2009)		-6.0%

^aOur empirical counterfactual does not disentangle the expenditure channel from the transfer channel (labeled Micro channels). For observations of older parents whose children have left the household, we impute some additional expenditures to parents of two children (Additional transfer channel) according to two scenarios (more or less conservative) based on the saving behavior of older parents still residing with their child.

Notes: Counterfactuals are run using estimates from the twins regressions. Macro (composition) channels are computed by multiplying the number of individuals born after 1980 by 2, at the same time imputing them lower incomes/education attainment as predicted by Table 8. Micro-channels are calculated using the response of expenditures of households at various ages of the children (for educ. exp.) and various ages of the parents (for non-educ. exp.) from Table E.1. See details in Appendix E.2. Model estimates in 2009 are based on running a ‘two-children’ policy in the quantitative model.

Model Counterfactuals. The rise in aggregate saving and human capital as predicted by the quantitative model can be viewed as an upper-bound of the overall effect of the one child policy (as

⁵⁰Our empirical counterfactual assumes that having twins is similar to having two children sequentially for saving decisions. A difference is that the arrival of twins may have been unanticipated, and another one is that there may be a difference in the degree of scale economies when having twins compared to having two children sequentially. The implications of this assumption for our results are discussed in Appendix E.2.

it assumes that the natural fertility has stayed constant). Ideally, one would like to know how much these variables would have increased in the absence of any fertility policies. Then, by comparing these outcomes with the one-child policy simulation results, one could attribute how much of the change could be tied to the policy itself. The challenge, though, is that one cannot observe variations in the data that would provide estimates of the natural fertility rate, and thus any estimate risks being speculative.

Nevertheless, one can still evaluate the overall effect of the policy in this model under different hypotheses for the path of natural fertility. A first approach, similar to our empirical counterfactual, is to assume that over the period considered, the natural fertility rate of China would have stayed above 2. In this case, a ‘two-children policy’ implemented post-1978 provides a lower-bound for the overall effect of the policy. A second approach is to assess the natural fertility rate in China over the period based on a fertility-income relationship observed in a cross-section of countries. We follow these two approaches sequentially. Details of these counterfactuals together with outcomes of the simulations are relegated to Appendix D.

‘Two-children’ policy. In line with the two children limit implemented in 1978, we implement a ‘two-children policy’ by assuming that fertility declines progressively over the period 1970-1977 before reaching the limit — $n_{\max,t} = 1$ for $t \geq 1978$.⁵¹ All other parameters of the model are set to their baseline value of Table 3. Under such a policy, the quantitative model predicts a 6.2% lower aggregate saving rate in 2014 than that under the one-child policy—about a third of the increase in aggregate saving rate over the last thirty years. Note that the prediction on aggregate saving is in the ballpark of our empirical estimates (Table 10). The human capital of the generation born in the mid-1980s is predicted to be 24% lower than under the one-child policy. We view these numbers as being conservative lower-bounds of the overall effect as fertility falls to 2 very early on in this simulation.

Natural fertility rate. We now investigate a counterfactual path of fertility in the absence of any fertility constraints. So far, with a constant preference for fertility v , equal to its pre-1970 level, the counterfactual fertility rate for China remains also at its pre-1970 value — about 3 children. But given that China’s income has been rising rapidly since 1970, one may want to relax this assumption. The way we go about this is to take a short-cut in modelling the robust negative relationship between income and fertility observed in the data (Jones, Schoonbroodt and Tertilt (2010)) by assuming that, starting 1970, the preference for fertility v falls as income rises.⁵² We discipline the path of fertility preferences v_t to match the fertility-income relationship found in the data for a large cross-section of countries in 2000. More specifically, we compute the path of v_t such that, in equilibrium, the number of children $N_t = 2n_t$ born in a household at date t depends on the parental income $w_{\gamma_n,t}$ as follows:

$$N_t = \underline{N} + aw_{\gamma_n,t}^{-b} \quad (14)$$

where the asymptotic fertility rate \underline{N} and the parameters a and b are estimated in the cross-section of countries in 2000 — details of the estimation are provided in Appendix D. We then simulate our quantitative model assuming the path of v_t for which the fertility-income relationship of Eq. 14 holds — keeping all other parameters to their baseline value.⁵³ We find that the natural fertility rate falls

⁵¹We assume that fertility falls linearly in the early 1970s. Households starting to conceive before 1978 are also constrained by the limit implemented later on (see Section 2 and Appendix B).

⁵²We assume that the link between fertility and income is driven by preferences v , which depend on the level of income. A more sophisticated model linking fertility and income through— for instance— a higher opportunity cost of time raising children as income rises, is beyond the scope of our paper (see Jones, Schoonbroodt and Tertilt (2010)).

⁵³We also provide sensitivity analysis for the natural fertility rate around this baseline scenario: a scenario where the

progressively starting 1970 but at a much slower speed than under the one-child policy — fertility reaching 2 children per household in the early 2000s. The human capital of a generation born in 1985 is only 10% higher than their parents, compared to about 50% under the one-child policy. The rise in the aggregate saving rate over the period 1982-2014 is 5% compared to more than 11%—implying that the one-child policy would have contributed to 35% of the increase in the saving rate in the data.

6 Conclusion

We show in this paper that exogenous fertility restrictions in China may have led to a rise in human capital and in household saving rate—by altering saving decisions at the household level, and demographic and income compositions at the aggregate level. We explore the quantitative implications of these channels in a model linking fertility, human capital and saving through intergenerational transfers that depend on the quantity and quality of offspring. Saving predictions across ages also become distinct from that of the standard lifecycle model—where human capital investment and intergenerational transfers towards the elderly are absent. We show that where our quantitative framework can generate both a micro and macro effect on saving that is close to the data, the standard OLG model falls short on both fronts.

The impact of twins estimated from the data provides an out-of-sample check to our model predictions, based on a similar twin experiment. The impact on household saving, expenditures and the degree of the quantity-quality trade-off is very close between model and data estimates. We find that the ‘one-child policy’ can account for at least a third of the rise in the aggregate household saving rate since its enforcement in the early 1980s. Importantly, the micro-channel accounts for the majority of the effect. This contrasts with the standard lifecycle hypothesis which conventionally focuses only on the macro channel of shifting demographic compositions. The policy also significantly fostered human capital accumulation of the only child generation.

This paper demonstrates that shifts in demographics as understood through the lens of a lifecycle model remain to be a powerful factor in accounting for the high and rising national saving rate in China—when augmented with important features capturing the realities of its households. The tacit implication—on a broader scale—is that the one-child policy provides a natural experiment for understanding the link between fertility and saving behavior in many developing economies. The quantitative impact of the policy is still evolving as the generation of more-educated only children become older and exert a greater impact on the economy—both in human capital and demographic weight. We may therefore expect the effect of the policy on aggregate outcomes to remain in years to come, before the ageing of the generation of only child and the progressive relaxation of fertility constraints in China eventually reverse the effects.

asymptotic fertility rate \underline{N} is set to the replacement rate of 2 — significantly above our estimated baseline but within the 5% confidence interval; in a second scenario, we assume a constant elasticity to income corresponding to $\underline{N} = 0$. See Appendix D for details.

References

- [1] Allen, Franklin, Jun, Qian and Xian Gu. 2015. "China's Financial System: Growth and Risk", mimeo, Imperial College London.
- [2] Angrist, Joshua D., Lavy, Victor and Analia Schlosser. 2010. "Multiple experiments for the causal link between the quantity and quality of children." *Journal of Labor Economics* 28, 773-824.
- [3] Bai, Chong-en, Hsieh, Chang-tai and Yingyi Qian. 2006. "The Return to Capital in China." *Brookings Papers in Economic Activity*
- [4] Banerjee, Abhijit, Meng, Xin, Porzio, Tommaso and Nancy Qian. 2014. "Fertility Change and Household savings in China: Evidence from a General Equilibrium Model and Micro Data." Yale University working paper.
- [5] Banerjee, Abhijit, Meng, Xin and Nancy Qian. 2010. "The Life Cycle Model and Household savings: Micro Evidence from Urban China." Unpublished.
- [6] Barro, Robert J. and Gary S. Becker. 1989. "Fertility Choice in a Model of Economic Growth". *Econometrica* 57, 481-501.
- [7] Barro, Robert and Jong-Wha Lee. 2010. "A New Data Set of Educational Attainment in the World, 1950-2010." *Journal of Development Economics* 104(C), 184-198.
- [8] Becker, Gary, S. and H. Gregg Lewis. 1973. "On the Interaction between the Quantity and Quality of Children." *The Journal of Political Economy*, 81, 279-288.
- [9] Blanchard, Olivier J. and Francesco Giavazzi. 2005. "Rebalancing Growth in China: a Three-Handed Approach". CEPR Discussion Paper Series 5403.
- [10] Boldrin, Michele and Larry E. Jones. 2002. "Mortality, Fertility, and Saving in a Malthusian Economy." *Review of Economic Dynamics* 5, 775-814.
- [11] Browning, Martin. 1992. "Children and Household Economic Behavior." *Journal of Economic Literature* 30 (3), 1434-1475.
- [12] Bussiere, Matthieu, Yannick Kalantzis, Romain Lafarguette and Terry Sicular. 2013. "Understanding Household savings in China: the Role of the Housing Market and Borrowing Constraints." University Library of Munich MPRA Paper 44611.
- [13] Cai, Yong. 2010. "China's Below-Replacement Fertility: Government Policy or Socioeconomic Development?." *Population and Development Review* 36 (3), 419-440.
- [14] Carroll, Christopher D. and David N. Weil. 1994. "Saving and growth: A Reinterpretation," *Carnegie-Rochester Conference Series on Public Policy*.
- [15] Chakrabarti, Rupa. 1999. "Endogenous fertility and growth in a model with old age support." *Economic Theory* 13, 393-416.
- [16] Chamon, Marcos and Eswar Prasad. 2010. "Why Are Saving Rates of Urban Households in China Rising?." *American Economic Journal: Macroeconomics* 2(1), 93-130.
- [17] Chesher, Andrew. 1998. "Individual Demands from Household Aggregates: Time and Age Variation in the Composition of Diet." *Journal of Applied Econometrics* 13 (5), 505-524.
- [18] Cisno, Alessandro and Furio Rosati. 1996. "Jointly determined saving and fertility behavior: theory and estimates for Germany, Italy, UK, and USA." *European Economic Review* 40, 1561-1589.

- [19] Coeurdacier, Nicolas, Stephane Guibaud and Keyu Jin. 2015. "Credit Constraints and Growth in a Global Economy." *American Economic Review* 105 (9), 2838-2881.
- [20] Coeurdacier, Nicolas, Stephane Guibaud and Keyu Jin. 2014. "Fertility Policies and Social Security Reforms in China." *IMF Economic Review* 62(3), 371-408.
- [21] 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.
- [22] Du, Qingyuan and Shang-Jin Wei. 2013. "A Theory of the Competitive Saving Motive." *Journal of International Economics* 91 (2), 275-289.
- [23] Ge, Suqin, Dennis T. Yang and Jusen Zhang. 2012. "Population Policies, Demographic Structural Changes, and the Chinese Household Saving Puzzle ", IZA Discussion paper 7026.
- [24] Gourieroux, Christian S., Alain Monfort and Alain Trognon. 1984. "Pseudo Maximum Likelihood Methods: Applications to Poisson Models." *Econometrica* 52 (3), 701-720.
- [25] Guo, Si and Fabrizio Perri. 2013. "The allocation puzzle is not as bad as you think", mimeo, University of Minnesota.
- [26] Horioka, Charles Y. and Junmin Wan. 2007. "The Determinants of Household Saving in China: A Dynamic Panel Analysis of Provincial Data." *Journal of Money, Credit and Banking* 39, 2077-2096.
- [27] Horioka, Charles Y. 2014. "Why Do People Leave Bequests? For Love or Self-Interest? Evidence from a New International Survey of Bequest Plans." UP School of Economics Discussion Papers 201406, University of the Philippines School of Economics.
- [28] Huang, Wei, Xiaoyan Lei and Yaohui Zhao, 2016. "One-Child Policy and the Rise of Man-Made Twins." *The Review of Economics and Statistics* 2016 98:3, 467-476.
- [29] Imrohorglu, Ayse and Kai Zhao, 2017. "The Chinese Saving Rate: Long-Term Care Risks, Family Insurance, and Demographics." Mimeo USC Marshall School of Business.
- [30] Jones, Larry E., Alice Schoonbroodt and Michele Tertilt, 2010. "Fertility Theories: Can they explain the Negative Fertility-Income Relationship?," *Demography and the Economy*, 43-100, edited by John Shoven, University of Chicago Press.
- [31] Jones, Larry E. and Michele Tertilt, 2008. "An Economic History of Fertility in the U.S.: 1826-1960." *Frontiers of Family Economics*, Vol. 1 edited by Peter Rupert, Emerald Press,
- [32] Li, Hongbin, Jusen Zhang and Yi Zhu. 2008. "The Quantity-Quality Trade-Off of Children in a Developing Country: Identification Using Chinese Twins." *Demography* 45, 223-243.
- [33] Mankiw, N. Gregory, David Romer and David N. Weil. 1992. "A Contribution to the Empirics of Economic Growth", *Quarterly Journal of Economics*, 107, 407-437.
- [34] Manuelli Rodolfo E. and Ananth Seshadri. 2007. "Explaining International Fertility Differences." *Quarterly Journal of Economics* 124(2): 771-807.
- [35] Modigliani, Franco. 1986. "Life-cycle, individual thrift, and the wealth of nations'." *American Economic Review* 76(3), 297-313.
- [36] Modigliani, Franco and Shi Larry Cao. 2004. "The Chinese Saving Puzzle and the Life-Cycle Hypothesis." *Journal of Economic Literature* 42, 145-170.

- [37] Newman, John L. 1983. "Economic Analyses of the Spacing of Births." *American Economic Review* 73 (2), 33-37.
- [38] Li, Haizheng, Yunling Liang, Barbara Fraumeni, Zhiqiang Liu and Xiaojung Wang. 2013. "Human Capital in China, 1985-2008." *Review of Income and Wealth* 59(2), 212-234.
- [39] Oliveira, Jacqueline. 2016. "The Value of Children: Intergenerational Transfers, Fertility, and Human Capital." *Journal of Development Economics* 120, 1-16.
- [40] Qian, Nancy. 2013. "Quantity-Quality and the One Child Policy: The Only-Child Disadvantage in School Enrollment in Rural China." NBER Working Paper 14973.
- [41] Raut, Lakshmi K. and T.N. Srinivasan. 1994. "Dynamics of endogenous growth." *Economic Theory* 4, 777-790.
- [42] Rosenzweig, Mark R. and Kenneth I. Wolpin. 1980. "Testing the Quantity-Quality Fertility Model: The Use of Twins as a Natural Experiment." *Econometrica* 48, 227-240.
- [43] Rosenzweig, Mark R. and Kenneth I. Wolpin. 2000. "Natural "Natural Experiments" in Economics." *Journal of Economic Literature* 38 (4), 827-874.
- [44] Rosenzweig, Mark R. and Jusen Zhang. 2009. "Do population control policies induce more human capital investment? Twins, birthweight and China's 'one-child' policy" *The Review of Economic Studies* 76, 1149-1174.
- [45] Rosenzweig, Mark R. and Jusen Zhang. 2014. "Co-Residence, Life-Cycle savings and Intergenerational Support in Urban China." NBER Working Paper 20057.
- [46] Santos Silva, João and Silvana Tenreyro. 2006. "The Log of Gravity." *The Review of Economics and Statistics* 8, 641-658.
- [47] Scharping, Thomas. 2003. "Birth Control in China 1949-2000: Population Policy and Demographic Development." Routledge Curzon, London and New York.
- [48] Sianesi, Barbara and John van Reenen. 2003. "The Returns to Education: Macroeconomics. " *Journal of Economic Surveys* 17(2), 157-200.
- [49] Song, Zheng and Dennis T. Yang. 2010. "Life cycle earnings and the household saving puzzle in a fast growing economy." Unpublished.
- [50] Song, Zheng, Storesletten, Kjetil and Fabrizio Zilibotti. 2011. "Growing like China." *American Economic Review*, 101(1), 202-241.
- [51] Song, Zheng, Storesletten, Kjetil, Wang, Yikai and Fabrizio Zilibotti. 2015. "Sharing High Growth across Generations: Pensions and Demographic Transition in China." *American Economic Journal: Macroeconomics*, 7(2), 1-39.
- [52] Wang, Fei. 2012. "Family Planning Policy in China: Measurement and Impact on Fertility." University Library of Munich MPRA Paper 42226.
- [53] Wang, Xin and Yi Wen. 2012. "Housing Prices and the High Chinese Saving Rate Puzzle." FRB of St. Louis Working Paper No. 2012-038A.
- [54] Wan, Junmin. 2015. "Household saving and Housing Prices in China." *The World Economy* 38, 172-192.
- [55] Wen, Yi. 2011. "Making Sense of China's Excessive Foreign Reserves." Federal Reserve Bank of St. Louis Working Paper, 2011-006A.

- [56] Wei, Shang-Jin and Xiaobo Zhang. 2011. "The Competitive Saving Motive: Evidence from Rising Sex Ratios and saving Rates in China." *Journal of Political Economy*, 119 (3), 511-564.
- [57] Yang, Dennis T., Junsen Zhang and Shaojie Zhou. 2011. "Why are saving rates so high in China?" IZA discussion paper No. 5465.

A Data

A.1 Data Sources and Description — micro data

Common Definitions.

Nuclear household: a household with two parents (head of household and spouse) and either a singleton or twins.

Parents: a head of household and his or her spouse with at least one coresiding child.

Mother age at first birth: age of mother minus the age of the eldest child in the household.

Individual disposable income: annual total income net of tax payments: including salary, private business and property income, as well as private and public transfer income.

Household disposable income: sum of the individual disposable income of all the individuals living in the household.

Household saving rate: household disposable income less household expenditures as a share of household disposable income.

Individual saving rate: individual disposable income less individual expenditures as a share of individual disposable income.

1. Urban Household Survey (UHS)

We use annual data from the Urban Household Survey (UHS), conducted by the National Bureau of Statistics, for 1986 and 1992 to 2009. Households are expected to stay in the survey for 3 years and are chosen randomly based on several stratifications at the provincial, city, county, township, and neighborhood levels. Both income and expenditures data are collected based on daily records of all items purchased and income received for each day during a full year. No country other than China uses such comprehensive 12-month expenditure records. Households are required by Chinese law to participate in the survey and to respond truthfully, and the Chinese survey privacy law protects illegal rural residents in urban locations (Gruber (2012) ; Banerjee et al. (2014)).

The 1986 survey covers 47,221 individuals in 12,437 households across 31 provinces. For the 1992 to 2009 surveys the sample covers 112 prefectures across 9 representative provinces (Beijing, Liaoning, Zhejiang, Anhui, Hubei, Guangdong, Sichuan, Shaanxi and Gansu). The coverage has been extended over time from roughly 5,500 households in the 1992 to 2001 surveys to nearly 16,000 households in the 2002 to 2009 surveys.

Data preparation. We prepare the regression sample by performing several steps of data cleaning. The whole sample contains 195,227 household-year observations across 19 surveys between 1986 and 1992-2009. Some households (about 0.5% of all observations) for which the composition appears misreported are excluded. More precisely, households with either more than one head of household or more than one spouse of the head of households are excluded (0.32% of all observations). We also drop households with average age of parents below 18y or above 99y or with a mother age at first birth below 15 or above 55 (an additional 0.19% if observations). Finally, we drop 35 households with multiple simultaneous births of order higher than 3. Excluding these observations, the sample is made of 193,689 household-year observations. Finally, for our purpose, we restrict the sample to two-parent households with at least one co-residing child (140,009 household-year observations).

For Table 7, since the identification relies on a child leaving the household, we restrict the sample to households who appeared in more than one survey year with a coresiding child in at least one survey year between 1992 and 2009 (37,165 unique households appearing on average for 2.5 years).

Household composition. Unless stated otherwise, we limit the sample of households to those with children of 18 and below (or 21 and below) because older children who still remain in their parents' household are most likely income earners and make independent consumption decisions (rather than decisions being made by their parents). Starting with a sample of two-parent households with at least one co-residing child of 140,009 household-year observations, we end up with 85,891 (resp. 102,046) household-year observation with an eldest coresiding child up to 18y (21y). Children who have departed from their parents' household are no longer observed (unless they remain financially dependent). This may introduce a selection bias if a large number of children select into living independently based on (unobserved) characteristics correlated with our outcome variables. In practice, such selection is limited by the fact that: (i) children studying in another city are still recorded as members of their parents' household and (ii) less than 0.5% of surveyed individuals aged 18 to 21 years old are living without their parents in a uni-generational household.

When specified, we restrict our attention to nuclear households - households with an only child (or twins) and two parents. The sample of nuclear households with children up to 18y (resp. 21y) is made of 71,107 (83,067) household-year observations.

Twins. We identify a pair of twins as two children under the same household head who are born in the same year, and when available, in the same month. When comparing twins identified using year of birth data as opposed to using both year and month of birth data (available for 2007 to 2009), only 8 households out of 206 with children below 18 years were misidentified as having twins and only 1 nuclear household out of 154 was misidentified. Overall, twin households make up for roughly 1% of all households with young children, which is consistent with the biological rate of twins' occurrence.

Income and consumption.

Disposable Income. In the 1992-2009 UHS surveys, income is observed for each individual in the household. Disposable income is defined as the sum of salary, private business, property income and private and public transfers less income tax payments. For the year 1986, information on income is available only at the household level and is not disaggregated over the different income sources.

Labour Income. For UHS 1992-2009, labour income is defined as the sum of salary and business income (thus excluding property and transfers income). Real labour income is obtained by deflating the nominal figures by the nationwide urban CPI obtained from CEIC.

Household and individual consumption expenditures. Household consumption expenditures is the sum of the various components of household expenditures, including food, clothing, health, transportation and communication, education, housing (i.e. rent or estimated rent of owned house), and miscellaneous goods and services. Consumption data disaggregated across expenditure categories (and in particular the level of education expenditures) are only available for the years 2002 to 2006. Our definition of household consumption expenditures does not include interest and loan repayments, transfers and social security spending. Education transfers to children living in another city are only available for UHS 2002 to 2009 and, unless stated otherwise, are not included in the measure of consumption expenditures (exceptions being Table 5 columns 5 and 6, Table 6 columns 1 and 2 and Table 8).

Contrary to income, individual consumption expenditures are not directly observable. The empirical strategy developed in Coeurdacier, Guibaud and Jin (2015) and summarized in Appendix E.1 estimates age-specific individual consumption expenditures using household expenditures. When estimating individual consumption expenditures, we restrict our attention to individuals above 25 and income earners aged between 21-24 (with an annual income above 100 yuan). All individuals strictly

below 21 and those under 25 who do not qualify as income earners (unless they are the household head's spouse) are considered as children, whose consumption is thus imputed to other household members (typically their parents). For the year 1986, income is also not observed at the individual level and we use the same empirical strategy to estimate individual income using household income.

Educational Attainment. For all survey years between 1992 and 2009 we observe the highest level of education attainment for each individual in the household. Education attainments range from: (i) illiterate or semi-illiterate, (ii) primary school, (iii) lower middle school, (iv) middle level professional, technical or vocational school (i.e. technical secondary education), (v) upper middle school (i.e. academic secondary education), (vi) professional school (i.e. technical tertiary education) to (vii) college or above (i.e. academic tertiary education).

In Table 8, the following definitions apply for the education dummies:

Higher education: the dummy is equal to one if the child has reached post-secondary education (i.e. professional school or college or above).

Academic high school: the dummy is equal to one if the child's highest level of education is either an academic high school (upper middle school) or an undergraduate/postgraduate degree (college or above).

Technical high school: the dummy is equal to one if the child's highest level of education is either a technical or vocational high school or a professional school (i.e. junior college).

2. CHIP

The 2002 China Household Income Project survey provides detailed income and expenditures data for a sample of 6,835 urban households over 12 provinces. In the calibration of education expenditures, we use CHIP detailed education spending data. The following definitions apply:

Discretionary education expenditures: tuition and miscellaneous fees for non-compulsory education.

Compulsory education expenditures: sum of tuition and miscellaneous fees on compulsory education, expenditures on textbooks, boarding school fee and expenditures on nursery and kindergarten.

3. CHARLS

The China Health and Retirement Longitudinal Study (CHARLS) pilot survey was conducted in 2008 in two provinces—Zhejiang and Gansu. Subsequently, CHARLS conducted in 2011 is the first wave of the national baseline survey covering 28 provinces. Data for 2011 are now partially available. The main respondents are from a random sample of people over the age of 45, and their spouses. Detailed information are provided on their transfer received/given to each of their children. The urban sample in 2008 (2011) covers 670 households (4,224 households) of which 321 (1,699) have at least one parent above 60 and at least one adult children above 25.

Gross transfers: sum of regular financial transfer, non-regular financial transfer and non-monetary transfer (i.e. the monetary value of gifts, in-kind etc.) from adult children to elderly parents. In 2008, of the 359 urban households in which transfers occur between children and parents: regular monetary transfers represent 14% of the total value of transfer from children, non-regular monetary transfers represent 42%, and 44% takes in the form of non-monetary support.

Net transfers: gross transfers less the sum of all transfers from parents to children.

In Table D.2, we use CHARLS 2008 and focus on gross individual transfers from adult children towards their parents. We focus only on gross transfers because the Poisson estimation does not allow for negative values in the dependent variable. Note that negative net transfers between elderly parents and adult children occur in only 4% of the households in CHARLS 2008. The following definitions apply for Table D.2:

Transfers: the sum of all financial and non-monetary transfers from an individual child to his elderly parents.

Individual income: CHARLS 2008 does not provide data on children’s individual income. Therefore, in order to approximate the share of transfers in children’s income, we use UHS 2008 income data to predict the income of individuals in CHARLS 2008. We compute the average individual income level by province, gender and education level (four groups) for each 3-year age group in UHS. Then the incomes of these individuals with a certain set of characteristics are taken to be proxies for the incomes of children with the same set of characteristics in CHARLS. In CHARLS 2011 parents are asked to estimate each of their children’s household annual income. Regression estimates CHARLS 2011 using this measure are very similar.

Education level: categorical variable with 10 groups ranging from “no formal education” to “PhD level”.

5. Census

The 1982 Chinese census, accessed through IPUMS International, surveyed 1% of the Chinese population across 29 provinces. In the absence of data on registration (*hukou*) in the 1982 census, we classify households as urban if (i) no household members engage in agriculture and (ii) the household usual residence location is in a city or in a urban district. The urban sample includes more than 700,000 individuals.

The 1990 Chinese census surveyed 1% of the Chinese population across 31 provinces. As fertility restrictions are linked to one’s registration status at the time of the fertility decision, we define a household as urban if it satisfies three conditions: (i) it is currently residing in a city (ii) it has a registration status (*hukou*) for its current residence location (iii) it was already living in the same city in mid-1985. After excluding collective households (less than 0.2% of households), the urban sample includes 1.1 million individual observations.

In both censuses, the number of surviving children to a given household is known. However, a given child age is only reported if the child is still residing with the parents. As the first born child is likely to be the first to leave the parental household, we can infer the birth year of the first born only in households in which the number of coresiding children is equal to the number of surviving children. For an eldest child in the household of age up to 18y, only 7% of households in census 1990 (resp. 15% in census 1982) have more surviving children than coresiding children. At older ages of children, selection is more of an issue—only households who either had children at a later age or had children who stayed longer in the parental home can be used to compute the fertility by age of first-birth.

A.2 Data Sources and Description — macro data

1. Aggregate household saving

We use the CEIC China premium database for average disposable income and consumption expenditures time series. The underlying data are from the National Bureau of Statistics’ Urban Household Survey (UHS) and Rural Household Survey. Data are on an annual basis from 1980 to 2014. The advantage of CEIC relative to using directly the UHS micro-data is to provide a longer time-series.

Urban disposable income: average disposable income per capita (i.e. wage, household business and property income). Real urban disposable income is obtained by deflating nominal disposable income by the urban consumer price index from the same dataset.

Urban saving rate: disposable income per capita less consumption expenditures per capita as a share of disposable income per capita for urban households.

Rural saving rate: net income per capita less consumption expenditures per capita as a share of net income per capita for rural households.

2. Real interest rates and deposit to wealth ratio

Nominal and real deposit rate. Annual data on nominal deposit rates are from the People’s Bank of China (PBOC): 1 year and 5 years nominal deposit rates over the period 1979-2013. Data on annual inflation are from CEIC and National Bureau of Statistics (NBS). Annual real deposit rates are computed by subtracting annual inflation from the nominal deposit rate.

Marginal productivity of capital. Annual data on the return to capital in China over the period 1978-2014 are from Bai et al. (2006) — updated data compared to the published paper version. We use the return to capital in the non-agricultural sector as a baseline as we focus on urban households. Differences with their central estimate of the return to capital, including the agricultural sector, are very small.

Deposit to financial wealth ratio. We use various sources to compute a time-series of the ratio of deposit and cash to financial wealth in China. Data on balance-sheets of households for the year 1990 and 1996 are available from the NBS; years 1992-1997 are provided by PBOC and NBS, and the years 2004-2013 are from the CEIC. Our deposit to financial wealth ratio is the sum of deposit and currency holdings divided by the total financial wealth of households. We need to compute a series over the entire period of 1979-2013, used to measure the real interest rate faced by Chinese households (see Appendix D). For the years prior to 1990, we use the average over the years 1990 and 1992-1997 (we get similar results using the 1990 value). For 1991, we use the average between 1990 and 1992. For the years 1998-2003, we use the average between 1997 and 2004. The implied time series for the deposit to wealth ratio is shown in Figure D.2 in Appendix D. For comparison, Song et al. (2015) use a value of 81% — in the ballpark of the data we collected.

3. Fertility and GDP per capita across countries. Data are from the World Development Indicators (July 2016 version) for 181 countries in the year 2000. The data are used to estimate the relationship between fertility and income (see Section 5.2 and Appendix D). The measure of fertility is the total fertility rate (births per woman). It represents the number of children that would be born to a woman if she were to live to the end of her childbearing years and bear children in accordance with age-specific fertility rates in 2000. Income is measured by the GDP per capita in constant 2010 USD.

B Online Appendix: Policies and Fertility in urban China

Explaining the fall of fertility. Figure B.1 shows that fertility in urban China fell progressively for parents who started to conceive in the 1970s. In what follows, we endeavor to demonstrate that this progressive fall is largely the result of the fertility restrictions introduced in the late 1970s. To do so, we construct a counterfactual fertility path based on the following assumptions: parents having children in the 1970s are constrained by the 1978-1980 fertility limits (one-child policy) but exhibit, otherwise, the same fertility and birth-spacing patterns as parents with a first-born in 1964 (presumably unaffected by fertility policies). Our counterfactual exercise documents that the 1978-1980 policy can, alone, account for nearly all of the decline in the fertility of parents with a first birth in the 1970s. The earlier family planning policies adopted in the early 1970s (later, longer, fewer) seem to have more modest effects. Finally, we explore the possibility that a shift in preferences towards fewer children partially explains the fertility decline. Comparing the fertility behavior of Han Chinese (subject to the one-child policy) and ethnic minorities (subject to a two children limit) confirms that the one-child policy was, indeed, a binding constraint on fertility.




Figure B.1: Fertility by date of birth of the first child

Notes: Fertility corresponds to the number of surviving children by date of birth of the first-child. Data source: Census 1982 and 1990.

Data. Fertility patterns are calculated from both 1982 and 1990 census data. The advantage of the 1990 census is a larger sample size and a clearer urban/rural distinction provided by the hukou registrations. The number of children of a given family is known, however, the year of birth is only observed if children still reside in the family — implying that one cannot study fertility behaviour from much earlier than 1970. For this reason, the 1982 census is the primary source used to study earlier

fertility behavior, though the urban/rural distinction is less explicit (see Appendix A). Robustness checks performed using the 1990 census gave very similar results.⁵⁴

Pre-policy fertility patterns. We first investigate the fertility patterns and birth spacing for households whose first child was born in 1964 — when fertility was unconstrained and fertility policies were still quite far into the future. Figure B.2 shows the fraction of households (with a first-born in 1964) having their n -th child after a particular number of years. Households with multiple children typically have the second child after almost 3 years, the third after 5-6 years and the fourth after 8 years. This pattern is taken to be the baseline fertility behavior before the implementation of fertility policies.

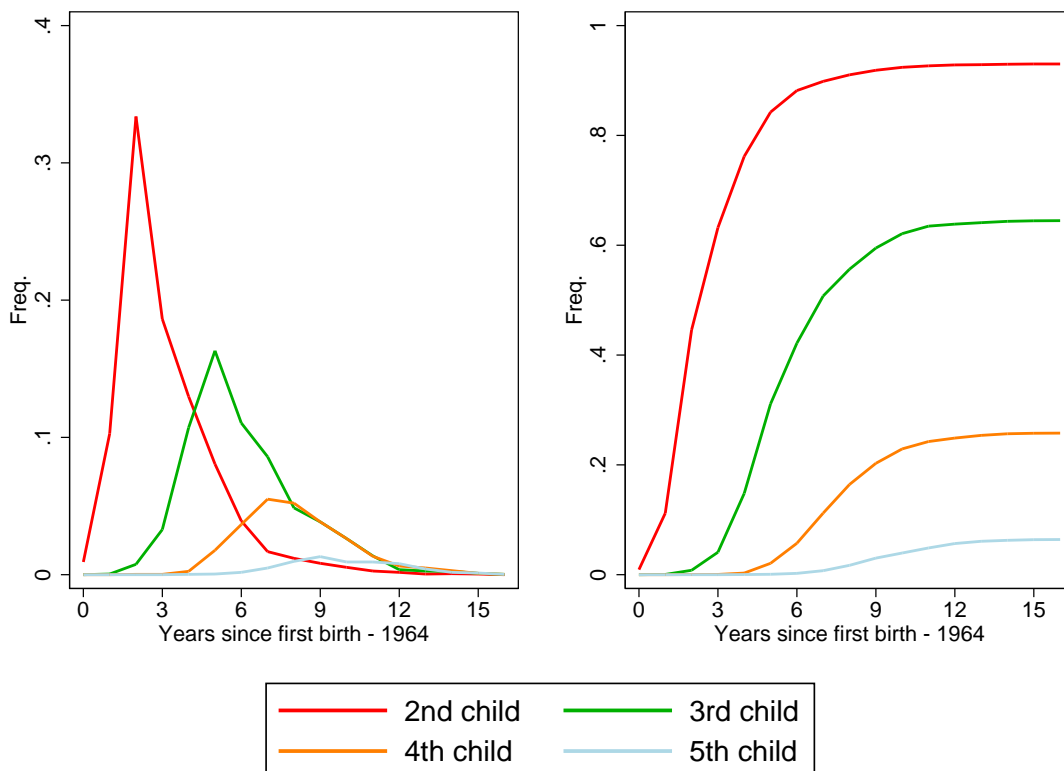


Figure B.2: Fertility patterns and birth-spacing in 1964.

Notes: The left panel shows the fraction of households with first-birth in 1964 having the n -th child after x years. The right panel shows the fraction of households with first-birth in 1964 having the n -th child before x years (cumulative distribution $D_n(x)$). Data source: Census 1982.

One-child policy counterfactual. Conceiving multiple children requires several years. As a consequence, parents with a first-born in the 1970s might have missed the window to bear additional children before the implementation of fertility restrictions in 1978-1980. To assess the quantitative importance of this mechanism, we assume that households bearing children post-1964 had the same fertility and birth spacing patterns as parents with a first child born in 1964, but became constrained by a two children limit in 1978 and a single child limit after 1980. Based on these assumptions, we compute the counterfactual fertility of urban Chinese households whose first child is born between

⁵⁴If we go further back with the Census 1990, the sample may suffer from selection as we only observe the fertility of households who either had children at a later age or had children who stayed longer in the parental home. Figure B.1 compares urban fertility (by date of first-birth) from both censuses. While both censuses exhibit a comparable fall in fertility over the period, the fertility measured from the 1990 census is slightly lower — particularly so in the late 1960s. This discrepancy is most likely driven by the selection of households with a first child born in the late 1960s in the 1990 census. Thus, in order to study pre-policy fertility behavior (in the 1960s), the baseline case takes the fertility patterns reported by the 1982 census.

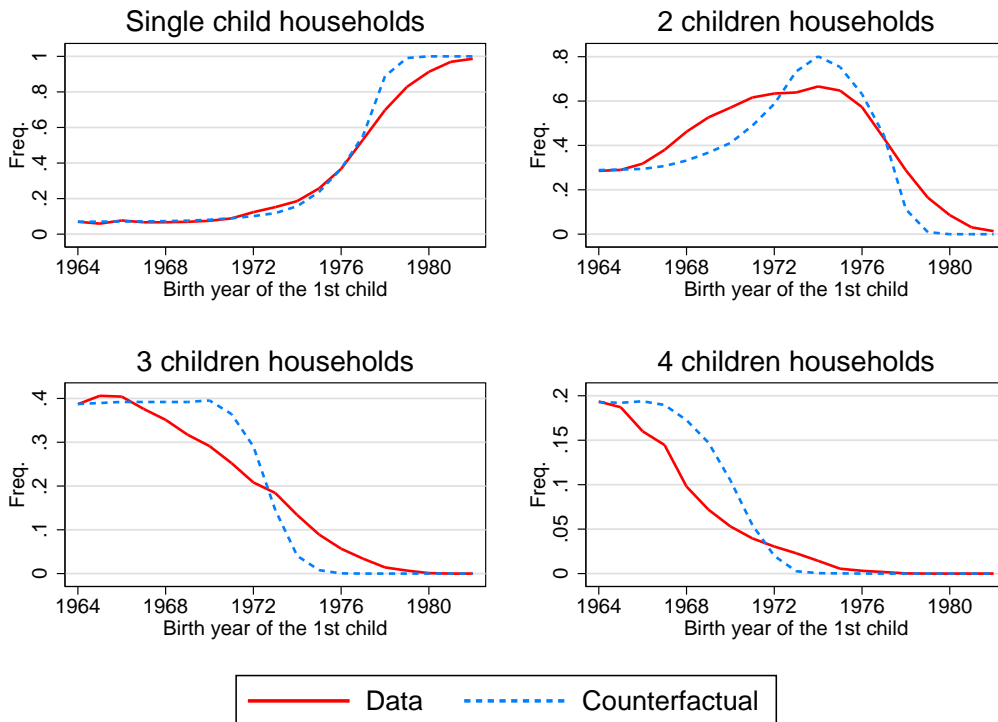



Figure B.3: Fertility (upper panel) and distribution of households by number of children (lower panel), by date of first-birth: Counterfactual versus Data.

Notes: For each date of first-birth, the upper panel shows the number of children in a household in our counterfactual and in the data. For each date of first-birth, the lower panel shows the distribution of households by number of children in our counterfactual and in the data, i.e. the fraction of households with n children. In our counterfactual, a two children limit is binding starting 1978 and a single child limit is binding starting 1980. Data source: Census 1982.

1964 and 1980. To understand our one-child policy counterfactual, consider households with a first-born in 1974. Due to the single child limit in 1980, the fraction of those households having a second child is set to the fraction of households having a second child in less than 5 years in 1964. Due to the two children limit in 1978, the fraction of them having n children, for $n > 2$, is the fraction having n children in less than 3 years in 1964. Our counterfactual fertility is computed by applying this reasoning to all households having a first-born between 1964 and 1980.⁵⁵

Our counterfactual, shown in Figure B.3 (upper panel), confirms that the 1978-1980 fertility restrictions can, alone, account for almost all of the progressive decline in the fertility of parents with a first-born in the 1970s.

Our counterfactual exercise also predicts the distribution of households by number of children and by date of first-birth (Figure B.3, lower panel). The counterfactual distribution fits the data reasonably well, even though it slightly overestimates the fraction of three and four children households in the first half of the 1970s. This could be explained by the conservative nature of our counterfactual exercise, assuming no fertility restriction before 1978 even though some provinces started implementing two children limits in the mid-1970s (see discussion below).

‘Wan, xi, shao’ policy. The ‘wan, xi, shao’ (later, longer, fewer) policies introduced in 1971 encouraged households to postpone the age of marriage, increase birth spacing, and bear fewer children—with a recommendation for 2 children only (Cai (2010) and Scharping (2003)). The timing and the extent of enforcement of these policies initially varied across provinces, but were gradually more uniform and stricter over the course of the decade. Overall the effect of these policies on fertility appears to be modest compared to the strict fertility limits enforced in the late 1970s.

Later. Data reveals that women did postpone their marriages—and consequently delayed their age of first-birth (see Figure B.4). On average, we find that women postponed the first-birth by an average of 28 months over the period 1970-1980. By 1985, average mother’s age at first-birth dropped significantly with the end of ‘wan, xi, shao’ policy—suggesting that the postponing of the first-birth was largely driven by the policy. Postponing the first-birth has an impact on fertility since older mothers have fewer children. Mothers aged below 24 at first birth in 1964 (‘younger’ mothers) have on average 3.1 children while those aged 24 or above (‘older’ mothers) have 2.6 children on average. Following the same strategy as before, we assume that the fertility difference between ‘older’ and ‘younger’ mothers did not change after 1964 to compute the counterfactual fertility due to later marriages only. We find that raising the share of ‘older’ mothers at first-birth from 45% in 1964 to 91% in 1981 (in line with Figure B.4) reduces completed fertility by 0.23 child—a very modest effect.⁵⁶

Longer. We provide evidence that birth spacing was barely affected by the fertility policies in the 1970s. To show this, we focus on households having a second child (resp. a third child), and compute the (conditional) distribution of birth-spacing between the first and second child at different dates of first-birth (resp. between the second and the third at different dates of birth of the second child).

⁵⁵In practice, we use the cumulative distribution function $D_n(x)$ which is the fraction of parents with a first child born in 1964 having a n -th child by date $1964 + x$ (see Figure B.2 (right panel) for a representation of $D_n(x)$). Due to the single child limit in 1980, the fraction of parents with a first-born at t (between 1964-1980) and having a second child is $D_2(1979 - t)$. Similarly, because of the two children limit in 1978, the fraction of parents having a n -th child ($n > 2$) is equal to $D_n(1977 - t)$. For example, parents with a first child born in 1964 and in 1974 have in the counterfactual the same probability $D_2(5)$ of having a second child five years after the first one. However, in 1980 those with a first child in 1974 can no longer have additional children.

⁵⁶This estimate is an upper bound of the effect of marriage postponement on realized fertility since the introduction of the one-child policy limited the ability of later marriage age to durably affect fertility.




Figure B.4: Share of ‘older’ mothers (24 and above) by date of first-birth.

Notes: For each date of first-birth, the plot shows the fraction of mothers aged 24 and above in both censuses. Data source: Census 1982 and Census 1990. Lower panel uses Census 1990.

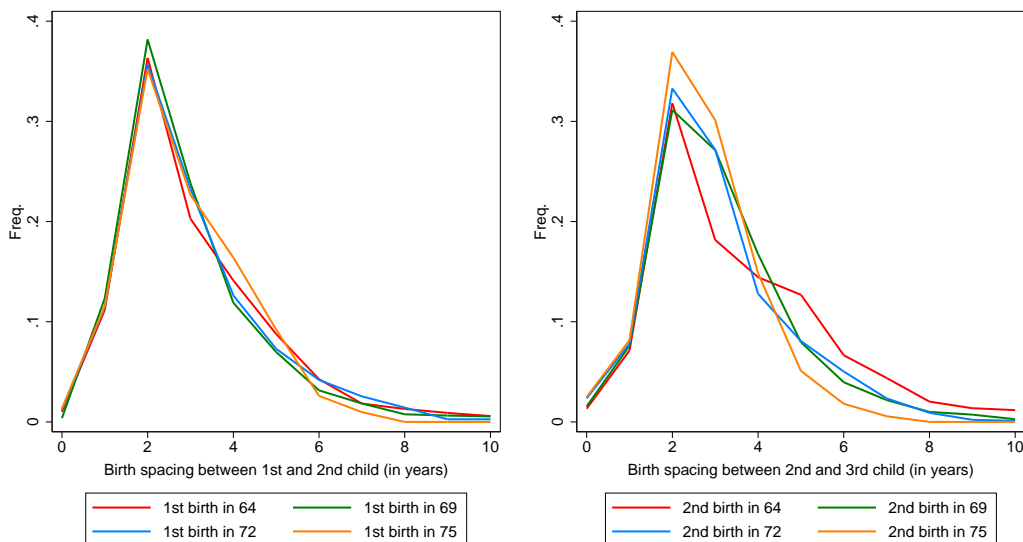


Figure B.5: Distribution of birth-spacing between the first and second child by date of birth of the first child (left panel) and between the second and third child by date of birth of the second child (right panel).

Notes: By date of first-birth (resp. second-birth), the distributions are conditional on having a second child (resp. a third child). We restrict our analysis to the period 1964-1975 since later on the birth of a second and/or third child is strongly affected by constraints on the number of children. Data source: Census 1982.

While the policy encouraged birth-spacing, we find that the distribution of birth-spacing between the first and the second child remains the same over the period 1964-1975 (Figure B.5, left panel). The distribution of birth-spacing between the second and the third-child hardly changed before and after the policy (1969 versus 1972). A difference is slightly more noticeable when comparing with 1975 but this difference is most likely driven by the constraints on the number of children implemented in the late 1970s.⁵⁷

Fewer. The government strongly encouraged households to have fewer children in the 1970s and ideally not more than two. A fertility limit at two children was initially introduced in some regions in 1973-74 and became a nation-wide policy with stricter enforcements in 1978 (Scharping (2003), p.51). The decline in the share of third and fourth order births which started in the mid-1970s and accelerated around 1978 supports this narrative (Figure 1 in the main text (upper panel)).

To sum-up, these results strongly indicate that, in line with our modeling assumption, quantitative limits on fertility are the main driving force behind the decline in fertility in urban China. More specifically, the 1978-1980 nation-wide quantity restrictions were the main source of fertility change. In contrast, other policies targeted toward marriage postponement and longer birth spacing seem to have had much more modest effects on fertility.

Preference for lower fertility? One could still argue that, instead of government policies, a shift in preferences towards fewer children partially explains the decline in fertility. To examine this possibility, we compare the fertility behavior of different ethnic groups which had different fertility restrictions. The non-Han minorities were not imposed a single child limit and were only constrained by a two children limit, enforced starting 1982 (Wang (2012)).

While both groups had roughly 3 children initially, the fertility rate fell to 1 for the Hans and to 2 for the non-Hans (Figure B.6, top panel) — consistent with their respective constraints. The role of the fertility policies, different for the two groups, is born out when inspecting the lower panel of Figure B.6, which shows the number of children in the years following the birth of the first child, by date of birth of the first child. While both groups displayed very similar fertility patterns prior to 1970,⁵⁸ they start to diverge in the 1970s: non-Han minorities have significantly more children in the long-run than Hans. The pattern is particularly striking for households with a first-born in 1975: while Han households had on average a bit more than 1.5 children by 1980, and no additional children thereafter, the non-Han minorities had reached a bit above 2 children ten years after the birth of the first child.

⁵⁷We focus on the evolution of birth-spacing over the period 1964-1975 since, later, the birth of a second and/or third child is strongly affected by constraints on the number of children. The distribution between the second and third child for a second birth in 1975 is already likely to be affected by the 1978 two children limit. Note also that we do not find evidence in our data that families reduced birth-spacing in anticipation of the one-child policy.

⁵⁸More precisely, before the two children limit implemented on Hans in the late 1970s, the increase over time in the size of a household with a first-born in 1969 is very similar for both ethnic groups.

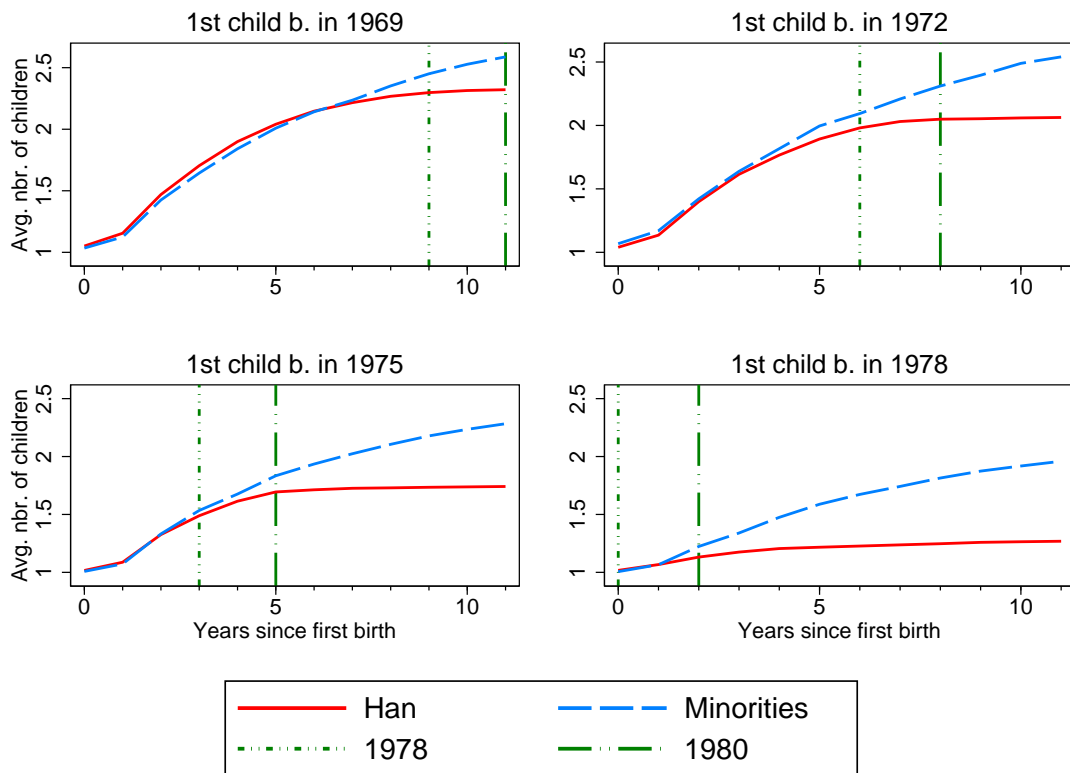



Figure B.6: Aggregate Fertility for Han and non-Han Minorities (upper panel) and number of children in Han and non-Han Minorities years after the first-birth in the household (lower panel).

Notes: Upper panel: Number of (surviving) children by date of birth of the first-born in the household. Lower panel: Average number of children in Han and non-Han Minorities households after a certain number of years following the first-birth in the household (for different dates of first-birth: 1969, 1972, 1975, 1978). Dashed vertical lines correspond to the two children limit in 1978 and the one-child policy in 1980. Data Source: Census 1990.

C Online Appendix: Theory

Proof of Proposition 1:

Proof of existence and uniqueness: if $\{n_{ss}; h_{ss}\}$ exists, then it must satisfy the steady-state system of equations:

$$\begin{aligned}\frac{n_{ss}}{1 - \theta - \psi \frac{n_{ss}^{\omega-1}}{\omega}} &= \left(\frac{v}{\beta(1 + \beta) + v} \right) \left(\frac{1}{\phi_0 + \phi_h (1 - \lambda) h_{ss}} \right) \\ h_{ss} &= \left(\frac{\alpha \psi \mu}{\phi_h} \right) \frac{n_{ss}^{\omega-1}}{\omega},\end{aligned}$$

which, combined, yields:

$$\frac{n_{ss}}{1 - \theta - \psi \frac{n_{ss}^{\omega-1}}{\omega}} = \left(\frac{v}{\beta(1 + \beta) + v} \right) \left(\frac{1}{\phi_0 + \frac{\alpha}{\omega} (1 - \lambda) (\psi \mu) n_{ss}^{\omega-1}} \right).$$

Let $N_{ss} = n_{ss}^{\omega-1}$, and rewriting the above equation yields

$$N_{ss}^{-1/(1-\omega)} - \left(\frac{v}{\beta(1 + \beta) + v} \right) \left(\frac{1 - \theta - \frac{\psi}{\omega} N_{ss}}{\phi_0 + (1 - \lambda) \mu \frac{\alpha \psi}{\omega} N_{ss}} \right) = 0$$

Define the function $G(x) = x^{-1/(1-\omega)} - \left(\frac{v}{\beta(1+\beta)+v} \right) \left(\frac{1-\theta-\frac{\psi}{\omega}x}{\phi_0+(1-\lambda)\mu\frac{\alpha\psi}{\omega}x} \right)$ for $x > 0$. Then,

$$\lim_{x \rightarrow +\infty} G(x) = \left(\frac{v}{\beta(1+\beta)+v} \right) \frac{\psi/\omega}{(1-\lambda)(\mu\frac{\psi\alpha}{\omega})} < 0 \text{ if } \lambda > 1, \text{ and } \lim_{x \rightarrow 0^+} G(x) = +\infty$$

We have:

$$G'(x) = -\frac{x^{-\omega/(1-\omega)}}{1-\omega} + \frac{v\psi/\omega}{\beta(1+\beta)+v} \frac{\phi_0 + (1-\theta)(1-\lambda)\alpha\mu}{\left(\phi_0 + \left((1-\lambda)\mu\frac{\alpha\psi}{\omega}x \right)^2 \right)}.$$

Two cases are:

- Case (1): if $\phi_0 + (1 - \theta)(1 - \lambda)\alpha\mu \leq 0$ then $G(x)$ is monotonically decreasing over $[0; +\infty]$.
- Case (2): $G(x)$ is first decreasing— to a minimum value strictly negative attained at $x_{\min} > 0$ — and then increasing for $x > x_{\min}$.

In both cases, the intermediate value theorem applies, and there is a unique $N_{ss} > 0$ such that $G(N_{ss}) = 0$ —thus pinning down a unique $\{n_{ss}; h_{ss}\}$ such that both are greater than 0. Moreover, if we define a unique n_0 implicitly by

$$\frac{n_0}{1 - \theta - \psi \frac{n_0^{\omega-1}}{\omega}} = \left(\frac{v}{\beta(1 + \beta) + v} \right) \left(\frac{1}{\phi_0} \right),$$

then it immediately follows that $n \geq n_0$ if $\omega \geq \alpha$ (and $\lambda > 1$).

Aggregate saving.

Definition of Saving Rates. The aggregate saving of the economy in period t , S_t , is the sum of the aggregate saving of each generation $\gamma = \{y, m, o\}$ coexisting in period t . Thus, $S_t \equiv \sum_{\gamma} S_{\gamma,t}$, where the overall saving of each generation $S_{\gamma,t}$ are by definition the change in asset holdings over a period with optimal asset holdings $a_{\gamma,t}$ given by Eq. 1 and Eq. 4: $S_{y,t} \equiv N_t^y a_{y,t}$, $S_{m,t} \equiv N_t^m (a_{m,t} - a_{y,t-1})$, and $S_{o,t} \equiv -N_t^o a_{m,t-1}$.

The individual saving rate $s_{\gamma,t}$ of cohort γ is the change in asset holdings over a period divided by the cohort's corresponding labor income (for the young and middle-aged) or capital income (for the old):⁵⁹

$$s_{y,t} \equiv \frac{a_{y,t}}{w_{y,t}}; \quad s_{m,t} \equiv \frac{a_{m,t} - a_{y,t-1}}{w_{m,t}}; \quad s_{o,t} \equiv -\frac{a_{m,t-1}}{(R-1)a_{m,t-1}} = -\left(\frac{1}{R-1}\right).$$

The aggregate saving rate, defined as $s_t \equiv S_t/Y_t$ (where Y_t denotes aggregate labor income), can thus be decomposed as follows:

$$s_t = s_{y,t} \left(\frac{n_t w_{y,t}}{n_t w_{y,t} + w_{m,t}} \right) + s_{m,t} \left(\frac{w_{m,t}}{n_t w_{y,t} + w_{m,t}} \right) + \frac{s_{o,t}}{n_{t-1}} \left(\frac{(R-1) a_{m,t-1}}{n_t w_{y,t} + w_{m,t}} \right). \quad (15)$$

The aggregate saving rate is thus a weighted average of the young, middle-aged and old's individual saving rates, where the weights depend on both the population and relative income of the contemporaneous generations coexisting in the economy—at a certain point in time. Changes in fertility can affect the aggregate saving rate through a *micro-economic channel*—changes in the individual saving behavior (change in $s_{m,t}$)—and a *macroeconomic channel*—changes to the composition of population and income.

Steady-State Aggregate saving. Long-run analysis helps gain intuition on how exogenous changes in long-run fertility impacts the aggregate saving rate. These exogenous changes can be brought about by a change in the preference for children ν , since it alters the birth rate but does not exert any impact on saving other than through its effect on n_{ss} . The saving rate, decomposed into the contribution of contemporaneous generations, is, in the long-run version of Eq. 15:

$$s = \frac{n_{ss} e}{(1 + n_{ss} e)} \underbrace{\left(\frac{-\theta \mu}{e} \right)}_{s_y} + \frac{1}{(1 + n_{ss} e)} \underbrace{\left(\kappa(n_{ss}) + \frac{\theta}{R} \right)}_{s_m} + \frac{\kappa(n_{ss})(R-1)}{n_{ss}(1 + n_{ss} e)(1 + g_z)} \underbrace{\left(\frac{-1}{R-1} \right)}_{s_o}, \quad (16)$$

where $\mu \equiv (1 + g)/R$, and $\kappa(n_t) \equiv a_{m,t}/w_{m,t}$ is given by the steady-state equivalent of Eq. 4:

$$\kappa(n_{ss}) = \frac{\beta}{1 + \beta} \left[(1 - \theta) - \underbrace{\left(\phi_0 n_{ss} + \alpha \psi \mu \frac{n_{ss}^\omega}{\omega} \right)}_{\text{cost of children} = \text{'expenditure channel'}} - \underbrace{\psi \frac{n_{ss}^{\omega-1}}{\omega}}_{\text{cost of parents}} - \underbrace{\frac{\psi \mu n_{ss}^\omega}{\beta \omega}}_{\text{benefits from children} = \text{'transfer channel'}} \right]$$

using $n_{ss} h_{ss} = \alpha \psi \mu n_{ss}^\omega / \omega$ from Eq. 6.

⁵⁹For analytical convenience, debt repayments for middle-aged and transfers are not included in the disposable income of the relevant generations. Results do not alter much except including more cumbersome expressions.

Proof of Lemma 1:

Substituting n_{\max} for the choice variable n_t in Eq. 6, the dynamics of $\log(h_t)$ becomes

$$\log(h_t) = \frac{1}{1-\alpha} \log\left(\frac{\alpha\psi}{\phi_h} \frac{n_{\max}^{\omega-1}}{\omega}\right) + \frac{1}{1-\alpha} \log(\mu_{t+1}) - \frac{\alpha}{1-\alpha} \log(h_{t-1}),$$

where $\log(h_t)$ is mean-reverting due to $-\frac{\alpha}{1-\alpha} < 1$ for $\alpha < 1/2$. It follows from $n_{t_0-1} > n_{\max}$ that $h_{\max} > h_{t_0-1}$. To assess the increase in human capital for the first generation of only child, we use we first use Eq. 6 to determine the human capital level in periods $t_0 - 1$ (in steady-state) and t_0 :

$$\begin{aligned} h_{t_0-1} &= \left(\frac{\alpha\psi}{\phi_h R} (1+g_z)\right) \frac{(n_{t_0-1})^{\omega-1}}{\omega} \\ (h_{t_0})^{1-\alpha} h_{t_0-1}^\alpha &= \left(\frac{\alpha\psi}{\phi_h R} (1+g_z)\right) \frac{(n_{\max})^{\omega-1}}{\omega} \\ \Rightarrow \left(\frac{h_{t_0}}{h_{t_0-1}}\right) &= \left(\frac{n_{t_0-1}}{n_{\max}}\right)^{\frac{1-\omega}{1-\alpha}} \end{aligned} \quad (17)$$

Proof of Proposition 2:

Define aggregate labor income in the economy to be the sum of income of the young and middle-aged workers $Y_{t+1} = (1+n_t e) N_{m,t+1} w_{m,t+1}$. Population evolves according to $N_{m,t+1} = N_{y,t} = n_{t-1} N_{o,t+1}$, and analogously, $N_{y,t+1} = n_t N_{y,t} = n_t N_{m,t+1}$. Cohort-level saving at date $t+1$ are respectively:

$$\begin{aligned} S_{y,t+1} &\equiv N_{y,t+1} a_{y,t+1} = -\theta n_t N_{t+1}^m \frac{w_{m,t+2}}{R} \\ S_{m,t+1} &\equiv N_{m,t+1} (a_{m,t+1} - a_{y,t}) \\ &= N_{m,t+1} \left[\frac{\beta w_{m,t+1}}{1+\beta} \left(1 - \theta - n_t \phi(h_t) - \frac{\psi n_{t-1}^{\omega-1}}{\omega}\right) - \frac{w_{m,t+2}}{R(1+\beta)} \frac{\psi n_t^\omega}{\omega} + \theta \frac{w_{m,t+1}}{R} \right] \\ S_{o,t+1} &\equiv -N_{t+1}^o a_{m,t-1} = -\frac{N_{m,t+1}}{n_{t-1}} \left[\frac{\beta w_{m,t}}{1+\beta} \left(1 - \theta - n_{t-1} \phi(h_{t-1}) - \frac{\psi n_{t-2}^{\omega-1}}{\omega}\right) - \frac{w_{m,t+1}}{R(1+\beta)} \frac{\psi n_{t-1}^\omega}{\omega} \right] \end{aligned} \quad (18)$$

Let $S_{t+1} = \sum_\gamma S_{\gamma,t+1}$ (where $\gamma \in \{y, m, o\}$) be aggregate saving at $t+1$, denoted, then the aggregate saving rate $s_{t+1} = S_{t+1}/Y_{t+1}$ can be written as

$$s_{t+1} = \frac{1}{(1+en_t)} \left[\begin{aligned} &-\frac{\theta}{R} n_t \left(\frac{w_{m,t+2}}{w_{m,t+1}}\right) + \frac{\beta}{1+\beta} \left(1 - \theta - n_t \phi(h_t) - \frac{\psi n_{t-1}^{\omega-1}}{\omega}\right) - \frac{\psi}{R(1+\beta)} \frac{n_t^\omega}{\omega} \left(\frac{w_{m,t+2}}{w_{m,t+1}}\right) + \frac{\theta}{R} \\ &-\frac{\beta}{(1+\beta)(1+g_z)n_{t-1}} \left(1 - \theta - n_{t-1} \phi(h_{t-1}) - \frac{\psi n_{t-2}^{\omega-1}}{\omega}\right) + \frac{\psi}{R(1+\beta)} \frac{n_{t-1}^{\omega-1}}{\omega} \end{aligned} \right]. \quad (19)$$

The aggregate saving rate in $t_0 + 1$, after the policy implemented in t_0 , is obtained by replacing $t+1$ by $t_0 + 1$ in Eq. 19 and n_t by n_{\max} . Using the optimal relationship between fertility and human capital along the transition path: $\phi_h n_{\max} h_{t_0} = \left(\frac{\alpha\psi}{R} (1+g_z) \left(\frac{h_{t_0}}{h_{t_0-1}}\right)^\alpha\right) \frac{n_{\max}^\omega}{\omega} = \left(\frac{\alpha\psi}{R}\right) \frac{n_{\max}^\omega}{\omega} \left(\frac{w_{m,t+2}}{w_{m,t+1}}\right)$, we have

$$s_{t_0+1} = \frac{1}{(1+n_{\max}e)} \left[\begin{aligned} &\frac{\theta}{R} \left(1 - n_{\max} \frac{w_{m,t+2}}{w_{m,t+1}}\right) + \frac{\beta}{1+\beta} (1-\theta) \left(1 - \frac{1}{n_{t_0-1}(1+g_z)}\right) \\ &-\frac{\psi}{R(1+\beta)} \frac{n_{\max}^\omega}{\omega} \left(\frac{w_{m,t+2}}{w_{m,t+1}}\right) (1+\beta\alpha) - \frac{\beta}{1+\beta} \phi_0 \left(n_{\max} - \frac{1}{1+g_z}\right) \\ &+\frac{\psi}{R(1+\beta)} \frac{n_{t_0-1}^{\omega-1}}{\omega} (1+\beta\alpha) - \frac{\psi\beta}{1+\beta} \left(\frac{n_{t_0-1}^{\omega-1}}{\omega} - \frac{1}{n_{t_0-1}(1+g_z)} \frac{n_{t_0-2}^{\omega-1}}{\omega}\right) \end{aligned} \right]$$

The aggregate saving rate s_t in the initial period $t = t_0$ is the steady-state equivalent of the above

equation. In order to find the difference $s_{t_0+1} - s_{t_0}$ we first obtain, with some algebraic manipulation:

$$\begin{aligned}
s_{t_0+1} &= \left(1 + \frac{(n_{t_0-1} - n_{\max})e}{1 + n_{\max}e}\right) s_{t_0} \\
&= \frac{1}{1 + n_{\max}e} \left[\begin{aligned} &-\frac{\theta}{R} \left(n_{\max} \left(\frac{w_{m,t+2}}{w_{m,t+1}} \right) - n_{t_0-1} (1 + g_z) \right) \\ &-\frac{(1+\beta\alpha)\psi}{R(1+\beta)\omega} \left(n_{\max}^\omega \left(\frac{w_{m,t+2}}{w_{m,t+1}} \right) - (1 + g_z)n_{t_0-1}^\omega \right) - \frac{\beta}{1+\beta}\phi_0 (n_{\max} - n_{t_0-1}) \end{aligned} \right] \\
&= \frac{1}{1 + n_{\max}e} \left[\begin{aligned} &-\frac{\theta}{R} (1 + g_z) \left(n_{\max} \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha - n_{t_0-1} \right) \\ &-\frac{(1+\beta\alpha)\psi}{R(1+\beta)\omega} (1 + g_z) \left(n_{\max}^\omega \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha - n_{t_0-1}^\omega \right) - \frac{\beta}{1+\beta}\phi_0 (n_{\max} - n_{t_0-1}) \end{aligned} \right].
\end{aligned}$$

Rearranging,

$$\begin{aligned}
s_{t_0+1} - s_{t_0} &= \frac{(n_{t_0-1} - n_{\max})e}{1 + n_{\max}e} s_{t_0} + \frac{\theta\mu}{1 + n_{\max}e} \left(n_{t_0-1} - n_{\max} \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha \right) \\
&+ \frac{\beta}{(1 + \beta)(1 + n_{\max}e)} \left[\phi_0 (n_{t_0-1} - n_{\max}) + \left(\alpha + \frac{1}{\beta} \right) \frac{\psi\mu}{\omega} \left(n_{t_0-1}^\omega - n_{\max}^\omega \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha \right) \right],
\end{aligned}$$

where $\mu \equiv (1 + g_z)/R$. To prove that $s_{t_0+1} - s_{t_0} > 0$, we first use Eq. 17. This implies that if $n_{t_0-1} > n_{\max}$, then

$$\begin{aligned}
n_{t_0-1} - n_{\max} \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha &= n_{t_0-1} \left(1 - \left(\frac{n_{\max}}{n_{t_0-1}} \right)^{1 - \frac{\alpha(1-\omega)}{1-\alpha}} \right) > 0 \\
n_{t_0-1}^\omega - n_{\max}^\omega \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha &= n_{t_0-1}^\omega \left(1 - \left(\frac{n_{\max}}{n_{t_0-1}} \right)^{\frac{\omega-\alpha}{1-\alpha}} \right) > 0
\end{aligned}$$

if $\omega > \alpha$.

Identification through twins.

From Eq. 6, the per-capita human capital of the twins (denoted $h_{t_0}^{twin}$) must satisfy:

$$(h_{t_0}^{twin})^{1-\alpha} h_{t_0-1}^\alpha = \left(\frac{\alpha\psi}{\phi_h} \mu \right) \frac{(2n_{\max})^{\omega-1}}{\omega} < \left(\frac{\alpha\psi}{\phi_h} \mu \right) \frac{(n_{\max})^{\omega-1}}{\omega} = (h_{t_0})^{1-\alpha} h_{t_0-1}^\alpha.$$

This leads immediately to the first testable implication.

Proof of Lemma 2:

From 20, we have:

$$s_{m,t_0+1} - s_{m,t_0} = \Delta s_m = \frac{\beta}{(1 + \beta)} \left[\phi_0 (n_{t_0-1} - n_{\max}) + \frac{(1 + \beta\alpha)\psi(1 + g_z)}{R\beta\omega} \left(n_{t_0-1}^\omega - n_{\max}^\omega \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha \right) \right]$$

The saving rate for a middle-aged agent in period $t + 1$ is $s_{m,t+1} \equiv (a_{m,t+1} - a_{y,t})/w_{m,t+1}$. By Eq. 19, we have

$$s_{m,t_0+1} - s_{m,t_0+1}^{twin} = \frac{\beta}{1 + \beta} \left[\phi_0 n_{\max} + \frac{(1 + \alpha\beta)\psi(1 + g_z)}{R\beta\omega} n_{\max}^\omega \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha \left(2^{\frac{\omega-\alpha}{1-\alpha}} - 1 \right) \right].$$

The micro-channel on aggregate saving of moving from $n_{t_0-1} = 2n_{max}$ to n_{max} in t_0 is, using Eq. 17:


$$\begin{aligned}
\Delta s_m(2n_{max}) &= \frac{\beta}{1+\beta} \left[\phi_0 n_{max} + \frac{(1+\beta\alpha)}{R\beta} \frac{\psi(1+g_z)}{\omega} n_{max}^\omega \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha \left(2^\omega \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^{-\alpha} - 1 \right) \right] \\
&= \frac{\beta}{1+\beta} \left[\phi_0 n_{max} + \frac{(1+\beta\alpha)}{R\beta} \frac{\psi(1+g_z)}{\omega} n_{max}^\omega \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^\alpha \left(2^{\frac{\omega-\alpha}{1-\alpha}} - 1 \right) \right] \\
&= s_{m,t_0+1} - s_{m,t_0+1}^{twin}.
\end{aligned}$$

D Online Appendix: Quantitative Model

D.1 Timing of the model

An agent lives for 20 periods of 4 years. Between period 1 (age 1 to 4 years) and period 5 (age 17 to 20 years) the agent receives transfers from his/her parents, makes no independent consumption decision and receives no labor income. In period 6 (age 21 to 24 years), the agent enters the labor force and starts to receive labor income and make consumption decisions. In period 7 (age 25 to 28 years) he decides on the number and human capital level of his children and a new generation is born. The agent provides education transfers to children from period 7 (age 25 to 28 years) to period 14 (53 to 56 years) in line with the empirical evidence in Figure 7. In the end of period 15 (at age 60 years) the agent enters old-age, and from period 16 (age 61 to 64 years) to his death at the end of period 20 (age 77 to 80 years) he receives old-age support from his children. The timing of lifetime events is summarized in Figure D.1.

Figure D.1: Timing of Lifetime Events: Quantitative OLG Model



New agents are born every year between 1900 and 2150. The size of the cohorts born between 1900 and 1964 is set equal to the size of these cohorts in the 1982 Census (i.e. the number of observations of age 18 to 82 years in 1982). For generations born after 1964, the cohort size is derived endogenously from the model's fertility path. As discussed in Appendix B, the one-child policy also affected individuals conceiving in the 1970s and constrained to a single birth individuals having children after 1980. Depending on their year of birth, agents are affected to a different extent by fertility restrictions:

Not affected. Agents born strictly before 1946 make their fertility decisions before 1970 and are therefore not affected at all by fertility policies.

Partially affected. Agents born between 1946 and 1969 are partially affected by fertility policies. On the one hand, their human capital level and the number of siblings with whom they share the burden of supporting elderly parents are decided before any impact of fertility restrictions. On the other hand their decision on the number and human capital level of their children is constrained by fertility restrictions. Individuals born after 1956 are fully constrained by the single child limit and individuals born between 1946-1956 are also constrained but to a lower extent.

Fully affected. Agents born after 1970 are fully affected by fertility restrictions as both their parents' fertility decision and their own fertility decision are constrained by the policies. Only the individuals born (strictly) after 1980 are both only child and parents of only child.

D.2 Calibration

In this Section, we provide details on the calibration strategy in addition to the description in the main text. We follow the same order as in the main text. We focus on the parameters for which additional details compared to the main text are necessary.

Productivity growth. Our specification of the wage equation implies that productivity growth g can be estimated from the time trend of individual income:

$$w_{\gamma,t} = e_{\gamma}((1+g)^t z_0) H_t^{\alpha}$$

$$\log(w_{\gamma,t}) \simeq g * t + \log(z_0) + \log(e_{\gamma}) + \alpha \log(H_t)$$

Using individual level income data from 1992 to 2009 (UHS), we estimate productivity growth g by performing the following regression:

$$\log(w_{i,\gamma,h,p,t}) = cst + g * year_t + \alpha_{\gamma} + \alpha_p + \alpha_h + \alpha_{\gamma} + \varepsilon_{i,\gamma,h,p,t}$$

where $w_{i,\gamma,h,p,t}$ denotes the real salary and self-employment income⁶⁰ in year t of an adult i of age γ with h years of education and living in province p . α_p is a province fixed-effect. We control for the human capital (H_t) and the age-specific (e_{γ}) components by including years of education fixed effects (α_h) and age fixed effects (α_{γ}). Results are displayed in Table D.1. Annual productivity growth is estimated to be equal to 6.06% (Column (1)). As a robustness check we run a Poisson specification of Eq. D.2 to account for potential time trends in the extensive margin of employment (Column (2)). Results are barely affected. The baseline calibration uses the OLS estimate of Column (1) and we set $g = 6.1\%$ (annual basis).

Table D.1: Productivity growth

	(1)	(2)
	OLS	Poisson
Dependent variable	Log(earnings)	Earnings
Year	0.0606*** (0.000269)	0.0655*** (0.000339)
Additional Controls	YES	YES
Observations	302,429	359,820
(Pseudo) R-squared	0.360	0.413
Province Dummies	YES	YES

Notes: Data source: UHS (1992-2009). We take one observation per individual between the age of 21 and 60. Earnings is defined as the sum of salary and self-employment income. Additional controls include age dummies and years of education dummies. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.


Age income profile. Data on labour income across age-groups are used to calibrate the experience parameters $\{e_{\gamma}\}_{\gamma \geq \underline{\gamma}}$ where we normalize e_{12} to 1 (age [45-48]) (see Figure 6 in the main text). Labour income includes salary plus private business income. This gives the set of parameters for the vector $\{e_{\gamma}\}_{\gamma \geq \underline{\gamma}}$ shown in Table D.3.

Real interest rate. Our model relies on the assumption of a constant interest rate R faced by Chinese households. In particular, we abstract from general equilibrium effects through which fertility changes

⁶⁰Real salary and self-employment income are computed by deflating nominal salary and self-employment income from UHS by the nationwide urban CPI from CEIC. See Appendix A for details.

would affect the interest rate, which in turn would modify saving decisions. Such general equilibrium effects have been emphasized in Banerjee et al. (2014) and could potentially mitigate the impact of fertility on saving. We compute measures of interest rates faced by Chinese households. Due to substantial financial repression in China, Chinese households do not have full access to investment opportunities that offer rates close to the marginal product of capital (MPK) (see Allen et al. (2015) and Song et al. (2011, 2015)) : the majority of Chinese household saving is put into bank deposits (Figure D.2), and the deposit rate has largely been controlled and capped at an artificially low rate. In the data, the real rate of return on deposits decided by the government is much lower than the return to capital implied by the MPK as measured by Bai, Hsieh and Qian (2006): over the period 1979-2013, the average real 5 year deposit rate is 1.6% compared to 23.1% for the return to capital in the non-agricultural sector.⁶¹ In other words, the return on saving faced by households were largely determined by policy.

Figure D.2: Deposit to financial wealth ratio.



Notes: The deposit to financial wealth ratio is the ratio of deposit and cash to financial wealth. Data come from various sources detailed in Appendix A. Missing observations for which data have been interpolated are shown in dotted line (see Appendix A for details).

In the spirit of Curtis et al. (2015) (see also Song et al. (2015) for a similar approach), we assume that the rate of interest R_t faced by households is defined by:

$$R_t = \lambda_t R_t^d + (1 - \lambda_t) R_t^K \quad (20)$$

where R_t^d denotes the deposit rate which is controlled by the government and R_t^K denotes the return to capital implied by the marginal product of capital; λ_t measures the fraction of assets of households in the form of saving deposits (resp. $(1 - \lambda_t)$ measures the access of households to the MPK), which hovers between 70% and 90% in the data (Figure D.2).⁶²

This approach has one important advantage: R_t can be measured in the data — one can measure

⁶¹As our study focuses on urban households, we use the return to capital in the non-agricultural sector as a baseline. The aggregate return to capital in China measured by Bai et al. (2006) averages at 22.6% over the same period.

⁶² λ_t can also be interpreted as a measure the degree of financial repression.




Figure D.3: Real household interest rate.

Notes: The real interest rate faced by household is computed using Eq. 20. The dotted line indicates the average over the period 1979-2013. Data for the real deposit rate R_t^d , the return to capital R_t^K and the fraction of financial wealth held in deposits λ_t are detailed in Appendix A.




Figure D.4: Real household interest rate: sensitivity analysis.

Notes: The real interest rate faced by household is computed using Eq. 20. The solid line assumes a constant real deposit rate R_t^d equal to its average of 1.6% over the period 1979-2013. The dotted line assumes a constant a real deposit rate R_t^d equal to 1.6% and a constant deposit to wealth ratio λ_t equal to 80% (average over the period 1990-2013). Data for the real deposit rate R_t^d , the return to capital R_t^K and the fraction of financial wealth held in deposits λ_t are detailed in Appendix A.

R_t^d , R_t^K and the fraction of financial wealth held in deposits λ_t . It allows us to investigate if general equilibrium effects of fertility are large enough to show up when looking at the time series of R_t and if the assumption of a constant real rate is unreasonable for the period considered. We compute different time-series of R_t over the period 1979-2013 based on different assumptions for R_t^d and λ_t — the data used for the R_t^K are from Bai et al. (2006) and identical across time-series. Data used are described in Appendix A.

Our baseline time series for R_t shown in Figure D.3 uses the raw data for the real 5 year deposit rate and the deposit (and cash) to wealth ratio for λ_t .⁶³ Abstracting from some extreme variations driven by inflation, R_t is roughly constant over the period — averaging at 5.3%. This value is used in our calibration of the quantitative model. A roughly stable interest rate is the consequence of two forces: λ_t is slowly falling over time in the data due to a better access of households to financial markets (Figure D.2), which tends to increase R_t .⁶⁴ Simultaneously, the MPK is slightly decreasing over time which reduces R_t — an evolution potentially due to demographic changes as in Banerjee et al. (2014).⁶⁵ Keeping the same path for R_t^K , an alternative measure of R_t assumes a constant real deposit rate R_t^d equal to its average of 1.6% over the period (Figure D.4, solid line). Again, the real rate faced by households is fairly constant over the period. A last measure of R_t uses both a constant real deposit rate and a constant deposit to wealth ratio λ_t (Figure D.4, dotted line). In this latter case, the real rate is slightly falling in the most recent period due to the fall of the MPK.⁶⁶

Fertility, demographic structure and policy implementation. Given all other parameters, the preference for children parameter v is set to 0.58 in order to match the average fertility over the period 1964-1969 of 2.92 (Census 1982). Starting 1970, $n_{\max,t}$ vary to match the fertility observed in the data over the period 1970-1980. This illustrates the constraint imposed by the policy on households who started to conceive in the 1970s as discussed in Section 2, and detailed in Appendix B. For any date post-1980, fertility is constrained by the one-child policy: $n_{\max,t} = \frac{1}{2}$ for $t > 1980$. Figure D.5 shows the path of fertility in the model and in the data. As we do not have Census data for urban households prior to 1970, we set the initial population composition in 1964 such that it reproduces the size of each age group above 17 years old in the Census 1982. The size of the age groups above 60 in 1964 (bins 61-64, ..., 77-80) remains undetermined, as above 80 in 1982. Note that this is irrelevant for our purpose as they are not taking human capital decisions for the later cohorts and we focus on outcomes starting 1982 and they do not survive beyond this date. From this initial distribution, the population of each age groups evolves in line with the evolution of fertility in the model and the data (shown Figure D.5). The resulting population composition is shown in Figure D.6 together with the counterpart in the UHS data. Most of the evolution of the population composition is captured by our model but, due to the absence of mortality in our framework, we tend to overestimate the proportion of older individuals and respectively underestimate the proportion of middle-aged in the later years.

Estimation of ω and validation of the transfer function. We use CHARLS data to estimate


⁶³Data for the fraction of deposits in financial wealth do not cover the whole period considered. For the period 1978-1989, we use the average in the early nineties. For the missing observations starting 1990, we simply interpolate. Details are provided in Appendix A. Real deposit rate are computed using the nominal deposit rate net of CPI inflation. It exhibits some extreme variations driven by inflation — the reason why we also consider the case of a constant low real deposit rate of 1.6% — corresponding to the average over the period.

⁶⁴In the model of Song et al. (2011), households have increasing access to the marginal productivity of capital of private firms — corresponding to a fall in λ_t (see also Curtis et al. (2015), section IVD for a similar point).

⁶⁵The fall of R_t^K is more visible after 2007 and could be also partly linked to the 2008 financial crisis.

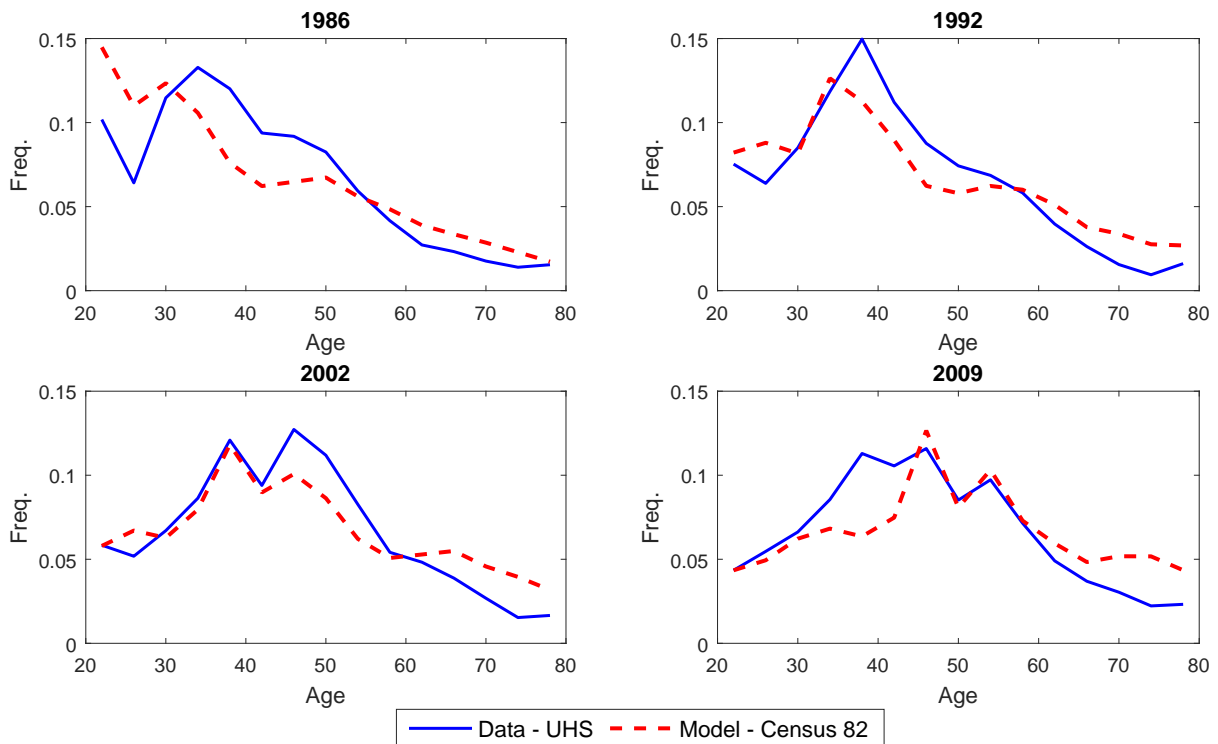
⁶⁶We also considered an alternative measure for the real deposit rate using 1 year nominal deposit rate. Results are not affected.

Figure D.5: Fertility. Model vs. Data



Notes: This figure shows the path of fertility constraints $n_{\max,t}$ imposed in our model (solid line) and the path of fertility in the Census data (dotted line). Fertility in the Census data measures the number of surviving children by average birth year of children in a household. See Appendix A for details.

Figure D.6: Population composition. Model vs. Data



Notes: This figure shows the share of each group in the population predicted by the model (solid line) and in the UHS data (dotted line).

the transfer function, $\psi \frac{n^{\omega-1}}{\omega} w$ (in logs). CHARLS provides data on transfers from a given child to his/her parents for the year 2008.⁶⁷ Using this cross-sectional data, the transfer function can be estimated by performing the following regression:

$$\log(T_{i,f,p}) = \alpha_p + \beta_n \log(n_f) + \beta_x \log(x_i) + \gamma Z_{i,f} + \varepsilon_{i,f,p}, \quad (21)$$

where $T_{i,f,p}$ denotes transfers per child i belonging to family f and living in province p to his/her parents. n_f denotes the number of children of a given family f , x_i a numerical indicator of quality of child i (education or imputed individual income),⁶⁸ $Z_{f,i}$ a vector of control variables (child's age and gender, child's and parents' age, dummy for the co-residence of parents) and α_p a province fixed-effect. The Poisson Pseudo-Maximum-Likelihood (PPML) estimator is employed to treat the zero values in our dependent variable (see Gourieroux, Monfort, and Trognon (1984) and Santos and Tenreiro (2006)).

Table D.2: Transfers from a given child to his/her parents

Dependent Variable	Transfers per child to parents		
	(1)	(2)	(3)
Log nbr. children	-0.349** (0.167)	-0.347** (0.171)	-0.330** (0.166)
Log educ. level	1.302*** (0.205)	1.207*** (0.191)	
Log income (predicted using UHS)			0.985*** (0.143)
Observations	1,489	1,489	1,475
Other controls	NO	YES	YES

Notes: Data source: CHARLS (2008). Sample restricted to children whose parents are above the age of 60. We take one observation per child. Estimation using Poisson Pseudo-Maximum-Likelihood (PPML). Robust standard errors in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Other controls included in all regression includes: age of child, average parents' age, a dummy for co-residence of the child with his parents and the child's gender.

Results are displayed in Table D.2. The amount of transfers (per offspring) given to parents is decreasing in the number of siblings the offspring has, and increasing in the offspring's quality — as either measured by education or income. The regression estimates for the elasticity of transfers to an offspring's income and to the number of his/her siblings correspond to our theoretical formulation of the transfer function, $\psi \frac{n^{\omega-1}}{\omega} w$ (in logs), with $\beta_n = \omega - 1$ and $\beta_x = 1$. The elasticity with respect to (imputed) income is very close to unity (Column 3), while the elasticity β_n of transfers to the number of children is equal to -0.35. Thus, ω is calibrated to 0.65.⁶⁹

Parameters $\{\beta, \psi, \theta\}$ and education parameters $\{\rho; \phi_\gamma; \phi_{\gamma,h}\}_{\gamma \in \{\gamma_n, \dots, \gamma_n + \gamma_e\}}$. Education expenditures observed in 2002 in CHIP can be decomposed between compulsory costs (tied to parameters

⁶⁷CHARLS include both rural and urban. We focus on urban households. When performing robustness checks on the whole sample of urban and rural, we find very similar results. We also perform robustness checks using the 'Three cities survey' for the year 1999 based only on urban households and the recent version of CHARLS (2011) with similar findings. See Appendix A for data description.

⁶⁸There is no direct income information for the children in CHARLS (2008). Therefore, we measure an offspring's quality x_i either by his/her education level (Columns 1-2); in Column 3, we use information on individual income and observable characteristics of the offspring (observed in UHS data) to assign to each child the income of an individual with the same set of characteristics in CHARLS data (see Appendix A).

⁶⁹In a non-reported regression using preliminary data from CHARLS (2011), we find a very similar estimate for ω ($= 0.61$) and a unitary elasticity w.r.t. income (CHARLS 2011 provides income data for the children). Using 'Three cities survey' data, we find a smaller estimate of ω ($= 0.52$) but not statistically different.

ϕ_γ) and discretionary costs (tied to parameters $\phi_{\gamma,h}$). Figure 3 provides details on this decomposition (see also Appendix A). Thus, computing compulsory expenditures as a share of wage income by age of parents directly pins down the values for the parameters ϕ_γ . We use the values shown in Table D.3 (values shown Fig. 7 in the main text).

While the parameters tied to compulsory education costs can be directly observed in the data, this is not the case of the parameters $\phi_{\gamma,h}$ tied to discretionary costs. Indeed, the model's counterpart of total discretionary education expenditures by age (in % of income) depends on the whole dynamics of human capital implied by the model and thus on all other parameters. However, if we assume that discretionary costs are zero up to age $\gamma = 8$ (age 29-32) — imposing $\phi_{\gamma,h} = 0$ for $\gamma \leq 8$, then one can solve analytically for education choices since they are not constrained by the borrowing limit.⁷⁰ More precisely, without binding borrowing constraints for education choices, the evolution of human capital satisfies under constrained fertility:

$$H_{t+\gamma_n} = \left(\frac{\kappa}{n_{t+\gamma_n}^{1-\omega}} \right)^{\frac{1-\rho}{1-(1-\rho)\alpha}} H_t^{\frac{\rho-\alpha(1-\rho)}{1-(1-\rho)\alpha}} \quad \text{with } \kappa = \left(\frac{(1-\rho)\alpha\psi \sum_{\gamma=\bar{\gamma}}^{\gamma_d} \left(\frac{1+g_z}{R}\right)^\gamma e_{\gamma-\gamma_n}}{\omega \sum_{\gamma=\gamma_n}^{\gamma_n+\gamma_e} \left(\frac{1+g_z}{R}\right)^\gamma \phi_{\gamma,h} e_\gamma} \right) \quad (22)$$

In this case, we also show that the model implied rate of change of discretionary education expenditures (in % of income) between two consecutive ages in the 2002 cross-section depends only on the dynamics of fertility which is known, the parameters α and ω which are calibrated independently, and the parameter ρ . To see this, denote $educ_{\gamma,2002}^m$ the share of income devoted to discretionary education expenditures at age γ in the model.⁷¹ For $\gamma > 8$, the rate of change between two consecutive ages is equal to:

$$\frac{educ_{\gamma+1,2002}^m}{educ_{\gamma,2002}^m} = \frac{\phi_{\gamma+1,h}}{\phi_{\gamma,h}} \cdot \frac{n_{2002+\gamma_n-\gamma-1}}{n_{2002+\gamma_n-\gamma}} \cdot \frac{h_{2002+\gamma_n-\gamma-1}}{h_{2002+\gamma_n-\gamma}} \quad (23)$$

All parents spending on education in 2002 were born under the pre-policy steady state and are therefore endowed with the same human capital level $H_{ss} = H_{t < 1970}$. Thus, absent binding credit constraint for education choices, one can solve analytically for the rate of change of human capital between two cohorts $\frac{h_{2002+\gamma_n-\gamma}}{h_{2002+\gamma_n-\gamma-1}}$ as a function of the rate of change of fertility between the same cohorts $\frac{n_{2002+\gamma_n-\gamma}}{n_{2002+\gamma_n-\gamma-1}}$ using Eq. 22:⁷²

$$\frac{h_{2002+\gamma_n-\gamma}}{h_{2002+\gamma_n-\gamma-1}} = \left(\frac{n_{2002+\gamma_n-\gamma}}{n_{2002+\gamma_n-\gamma-1}} \right)^{\frac{\omega-1}{1-\alpha(1-\rho)}} \quad (24)$$

From Eq. 23 and Eq. 24, we deduce :

$$\frac{\phi_{\gamma+1,h}}{\phi_{\gamma,h}} = \frac{educ_{\gamma+1,2002}^m}{educ_{\gamma,2002}^m} \cdot \left(\frac{n_{2002+\gamma_n-\gamma-1}}{n_{2002+\gamma_n-\gamma}} \right)^{\frac{\alpha(1-\rho)-\omega}{1-\alpha(1-\rho)}}$$

Thus, for a given ρ , there exists, up to a normalization, a unique vector of parameters $\phi_{\gamma,h}$ across ages which makes sure that: $\frac{educ_{\gamma+1,2002}^m}{educ_{\gamma,2002}^m} = \frac{educ_{\gamma+1,2002}^d}{educ_{\gamma,2002}^d}$, where $educ_{\gamma,2002}^d$ denotes the discretionary education expenditures (in % of income) at age $\gamma > 8$ in the 2002 data.⁷³ We pick the vector of $\phi_{\gamma,h}$

⁷⁰Note that this assumption is reasonable when looking at the data as discretionary education costs are almost zero when children are below 8.

⁷¹With the model's notations: $educ_{\gamma,2002}^m = \phi_{\gamma,h} h_{2002-\gamma+\gamma_n} n_{2002-\gamma+\gamma_n}$.

⁷²The human capital investment in a child born in $t = 2002 + \gamma_n - \gamma$ is equal to: $h_{2002+\gamma_n-\gamma} = \left(\frac{\kappa}{n_{2002+\gamma_n-\gamma}^{1-\omega} H_{ss}^{(1-\rho)\alpha}} \right)^{\frac{1}{1-\alpha(1-\rho)}}$

⁷³In the numerical exercise we set the first discretionary cost parameter $\phi_{9,h}$ equal to :

which satisfies the following equality for a given ρ :

$$\frac{\phi_{\gamma+1,h}}{\phi_{\gamma,h}} = \frac{educ_{\gamma+1,2002}^d}{educ_{\gamma,2002}^d} \cdot \left(\frac{n_{2002+\gamma_n-\gamma-1}}{n_{2002+\gamma_n-\gamma}} \right)^{\frac{\alpha(1-\rho)-\omega}{1-\alpha(1-\rho)}}$$

In other words, the parameters $\phi_{\gamma,h}$ are set to match the shape of the age-profile of discretionary education expenditures (in % of income) in the data.

Then, we search for the remaining parameters $\{\beta, \psi, \theta, \rho\}$ over a grid $\mathbf{\Gamma}$ in order to perform the following minimization between model's outcomes and data:

$$\min_{\{\beta, \psi, \theta, \rho\} \in \mathbf{\Gamma}} \left[\sum_{\gamma=\gamma}^{\gamma_d} \lambda_{\gamma}^s \left| s_{\gamma,1986}^m(\beta, \psi, \theta, \rho) - s_{\gamma,1986}^d \right| + \sum_{\gamma=\gamma_n}^{\gamma_m+\gamma_e} \lambda_{\gamma}^{educ} \left| educ_{\gamma,2002}^m(\beta, \psi, \theta, \rho) - educ_{\gamma,2002}^d \right| \right] \quad (25)$$

where $s_{\gamma,1986}^m$ (resp. $s_{\gamma,1986}^d$) is the model predicted saving rate at age γ in 1986 (resp. the saving rate at age γ in the 1986 data); λ_{γ}^s and λ_{γ}^{educ} are weights on different age groups summing to one and reflecting their respective income share (see Table D.3).

In practice, we use the following grid $\mathbf{\Gamma}$ with 12,540 unique combinations of the remaining parameters: β between 0.95 and 0.995 with a step size of 0.005; ψ between 4% and 13% with a step size of 0.5%, θ between 0 and 5% with a step size of 1% and ρ between 0 and 0.5 with a step size of 0.05. The minimization procedure leads to the values of parameters shown in Table 3: $\beta = 0.99$ (annual basis); $\psi = 9\%$; $\theta = 0\%$; $\rho = 0.2$. The corresponding discretionary education costs parameters $\{\phi_{\gamma,h}\}_{\gamma}$ are shown in Table D.3.⁷⁴

Table D.3: Age Dependent Parameters

Life-cycle															
Age group	21-24	25-28	29-32	33-36	37-40	41-44	45-48	49-52	53-56	57-60	61-64	65-68	69-72	73-76	77-80
Period γ	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Calibrated parameters															
e_{γ}	0.64	0.70	0.86	0.89	0.91	0.98	1.00	0.88	0.73	0.51	0.15	0.10	0.07	0.03	0.02
ϕ_{γ}	-	0.02	0.08	0.09	0.12	0.12	0.09	0.03	0.02	-	-	-	-	-	-
$\phi_{\gamma,h}$	-	-	0.00	0.01	0.01	0.04	0.07	0.04	0.02	-	-	-	-	-	-
Minimization weights															
λ_{γ}^s	0.05	0.06	0.10	0.13	0.14	0.11	0.12	0.11	0.07	0.04	0.03	0.02	0.010	0.01	0.00
λ_{γ}^e	-	-	0.15	0.15	0.16	0.15	0.15	0.13	0.11	-	-	-	-	-	-

Figure D.7 plots the shape of the objective function along the parameters β and ψ (fixing ρ and θ to their baseline values).⁷⁵ While we cannot formally prove the uniqueness of our combination of parameters, Figure D.7 is suggestive that for a fairly wide range of parameter values, the objective function exhibits a very noticeable minimum.


Lastly, we performed sensitivity analysis adopting different objective functions while keeping the same targets. In particular, instead of differences in absolute value, we used the squared values of the difference between model and data of saving rates across ages in 1986 and discretionary education

$\phi_{\gamma,h} = \left(\frac{n_{2002+\gamma_n-9}}{n_{ss}} \right)^{\frac{1-\omega}{1-\alpha(1-\rho)}} \frac{educ_{\gamma,2002}^d}{n_{2002+\gamma_n-9}}$. This normalization implies that the pre-one child policy steady state level of human capital is equal to one: $H_{ss} = \frac{\kappa}{n_{ss}^{1-\omega}} = 1$.

⁷⁴For computational efficiency reasons, we assume that after the credit constraint has bound a first time (for any number of consecutive periods) the agent is no longer financially constrained and can borrow below the credit limit. In practice, this assumption has no effects in most of the parameter space as the credit constraint does not generally bind more than once.

⁷⁵We focus on these two parameters as they are quantitatively the most relevant for our results. ρ is largely determined to match the levels of education spending across ages (second term in our minimization; see Eq. 25) and θ the saving rate of the young as shown in the sensitivity analysis.

Figure D.7: Objective function along parameters values.



Notes: The figure plots the log-value of the objective function shown in Eq. 25 along a grid of values for β and ψ . Other parameters are set to their baseline value (see Table 3).

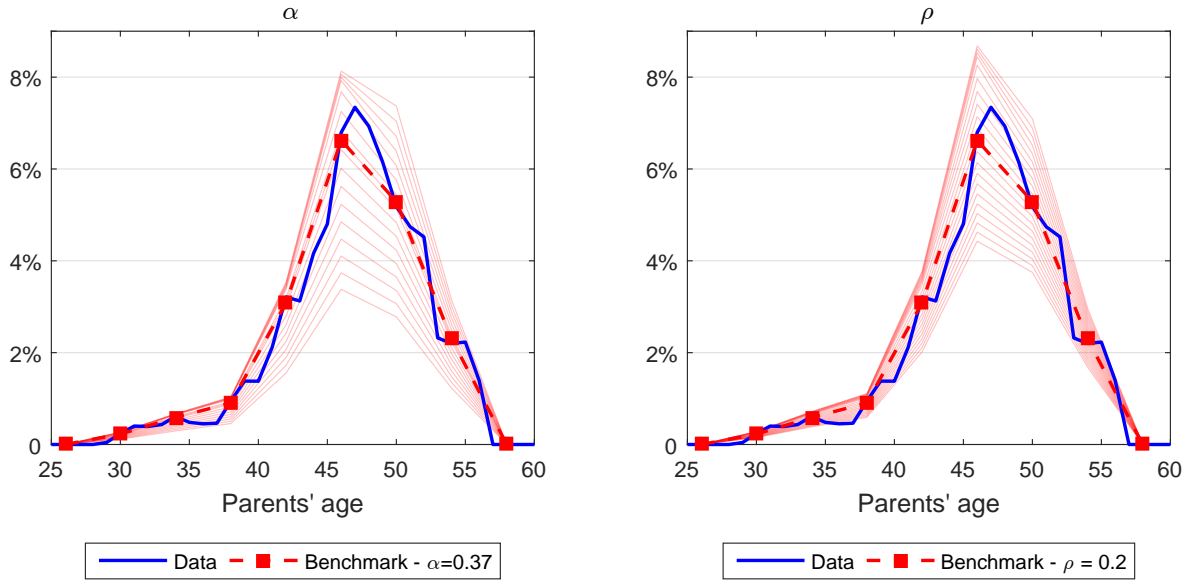
spending across ages (in % of income) in 2002. We also investigated alternative weighting schemes across age groups: equal-weighting or share of population-weighting. Values for the parameters were very close to our baseline calibration with these alternative objective functions and, most importantly, outcomes generated by the model were barely affected under these alternative calibrations.

D.3 Sensitivity analysis

In this Section, we provide sensitivity analysis around our baseline calibration for the main parameters of the model. This is informative regarding how parameters affect the shape of our model's outcomes — age-saving profiles and/or discretionary education expenditures. This also enlightens which parameters play a crucial role for our results. We provide sensitivity analysis along parameters that are not directly measured in the data: the measure of decreasing returns to education α , the degree of intergenerational transmission of human capital ρ , the discount rate β , the credit constraint parameter θ and the transfer parameter ψ . In our sensitivity analysis, we move one parameter along a set of possible values while keeping all other parameters to their baseline value (see Table 3). In our simulations, we find that the parameters α and ρ play an important role for education decisions and barely affect age-saving profiles while the parameters β , θ and ψ mostly affect life-cycle saving. For space considerations we only focus on outcomes which are significantly affected by a given parameter.

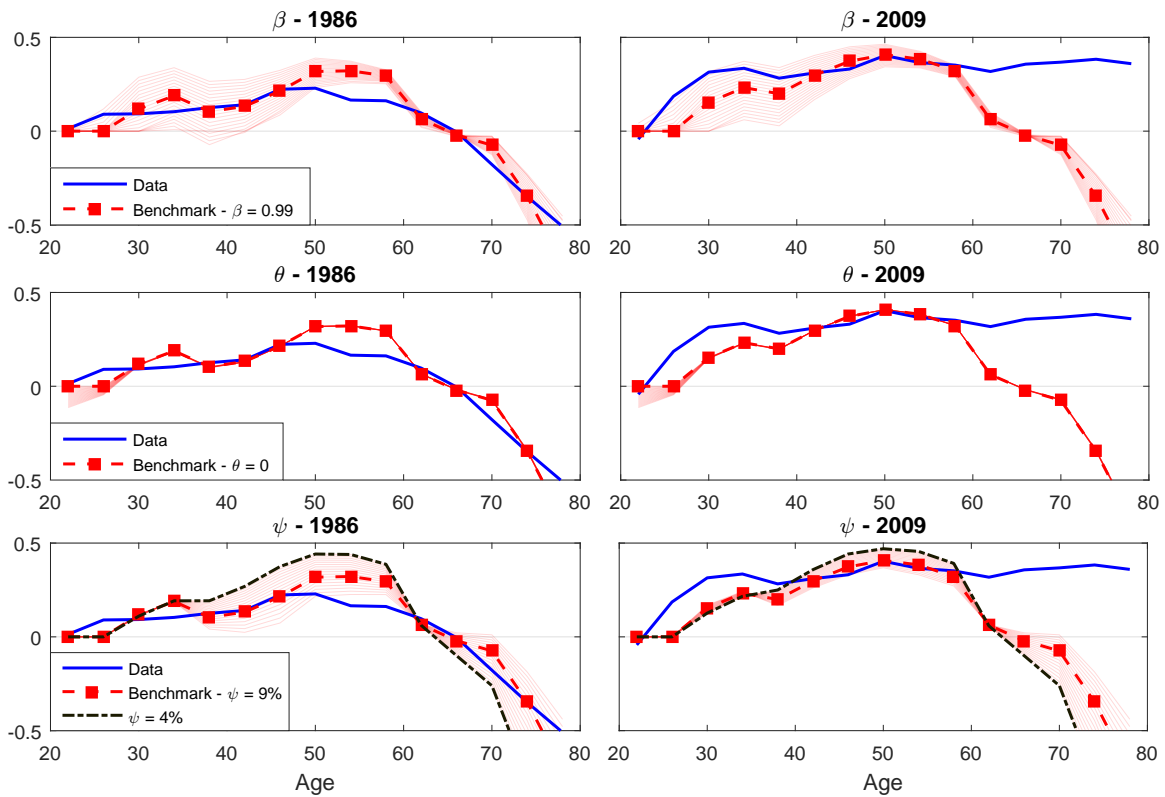
Sensitivity to parameters α and ρ . The curvature of technology with respect to human capital α and the degree of intergenerational transmission of human capital ρ are essential for matching the overall level of discretionary education spending in our calibration. This is shown in Figure D.8. A lower α or a higher ρ lowers the return to human capital investment and the level of discretionary education spending. It also makes the human capital response lower following the policy change. For instance, with $\alpha = 0.2$ (resp. $\rho = 0.4$), the increase in human capital of the generation of only-child compared to the generation of his/her parents (with two siblings) is 43% (resp. 34%) compared to 53% in our baseline. A lower α or a higher ρ would thus also imply a lower difference in discretionary education spending between parents of an only child and parents of twins—a moment our calibration tends to match relatively well. α and ρ barely affect age-saving profiles which are not shown.

Figure D.8: Sensitivity with respect to α and ρ : discretionary education spending (% of income).



Notes: These figures plot the model implied discretionary education spending (in % of income) in the 2002 cross-section for a range of values for α (left-panel) and ρ (right-panel); α is allowed to vary between 0.2 and 0.5, ρ between 0 and 0.4. When varying one parameter, all other parameters are kept fixed to their baseline value (see Table 3). Data from CHIP are shown for comparison purposes.

Figure D.9: Sensitivity with respect to β , θ and ψ : age-saving profile in 1986 (left panel) and 2009 (right panel).



Notes: This figures plots model implied the age-saving profile in 1986 (left panel) and 2009 (right panel) for a range of values for β (top-panel), θ (middle-panel) and ψ (bottom-panel); β is allowed to vary between 0.98 and 1, θ between 0 and 0.1 and ψ between 0.04 and 0.12. When varying one parameter, all other parameters are kept fixed to their baseline value (see Table 3). Data from UHS are shown for comparison purposes.

Sensitivity to parameters β , θ and ψ . These three parameters are relevant for generating age-saving profiles and aggregate saving in line with the data. An increase in β essentially increases saving at working age — except for the youngest ages due to the presence of credit constraints. As more wealth is accumulated before retirement with a higher β , dissaving at old age becomes also larger. Age-saving profiles in 1986 and 2009 for different values of β are shown in Figure D.9 (top panel). Thus, β plays an important role in matching the aggregate saving rate in the first years of the implementation of the one-child policy. As shown in Figure D.9 (middle panel), the parameter θ mostly affects the saving rate of young households which turn out to be credit constrained (age 21-28). While a low value of θ helps to match the saving rate of these households, it does not play much of a role quantitatively and our results remain valid for a wide range of values for θ . The transfer parameter ψ has a large impact on the shape of the age-saving profile particularly for those above age 40. It determines the magnitude of transfers received in old age and thus the need to save in middle-age. Figure D.9 (bottom panel) shows age-saving profiles for different values of ψ where the calibration with a low value of transfer ($\psi = 4\%$) is also shown for comparison purposes.

D.4 Natural fertility rate and counterfactuals

In this Section, we present further details regarding the counterfactuals where fertility in China is left unconstrained. Our strategy relies on feeding a path of fertility preferences v_t since 1970 such that our quantitative model reproduces the fertility-income relationship that can be observed in the data. As a short-cut, we embed a potential fall in the natural fertility rate through changes in preferences — a more sophisticated model linking fertility and income through, for instance, a higher opportunity cost of time raising children as income rises, being beyond the scope of our paper (see Jones, Schoonbroodt, and Tertilt (2010) for a survey).⁷⁶ We proceed in two-steps to build our counterfactuals: first, we estimate a fertility-income relationship in the cross-section of countries. Second, we set the preferences for fertility such that a simulation of the quantitative model, with all parameters but v_t set to their baseline, delivers a fertility-income relationship in line with the data.

Data and Estimation. We use cross-country data for the year 2000 to estimate the relationship between fertility and income per capita (data are described in Appendix A).⁷⁷ Based on the data shown in Figure D.10, we postulate the following parametric relationship between fertility and income:

$$N_i = \underline{N} + a(\text{GDP per capita}_i)^{-b} \quad (26)$$


where N_i denotes the fertility rate in country i (in 2000), GDP per capita_i the real GDP per capita in country i (in 2000). \underline{N} , a and b are positive parameters to be estimated. \underline{N} corresponds to an asymptotic fertility rate towards which countries would converge in the long-run as their income grows. As the non-linearity of fertility at high level of income seems an important feature of the data (see Figure D.10), our baseline specification relies on a non-linear least squares estimation (NLS), allowing us to estimate \underline{N} . For comparison purposes, using the same cross-section, we also estimate the following log-linear relationship —corresponding to $\underline{N} = 0$ (constant elasticity model) :

$$\log(N_i) = cste - b \log(\text{GDP per capita}_i) \quad (27)$$

⁷⁶Note that our theory generates endogenously part of the negative relationship between fertility and income due to the endogenous quantity-quality trade off. If fertility falls, income rises due to human capital accumulation. However, starting from a steady-state, our model does not generate any fall in fertility if income rises without shift in preferences.

⁷⁷Results are robust using the 2005 or 2010 cross-sections.

Figure D.10: Fertility and GDP per capita in 2000.



Notes: This figure plots the total fertility rate N_i as a function of GDP per capita $GDP\ per\ capita_i$ for a cross-section of countries in 2000. The red line shows the fit of the data as implied by a non-linear least square estimation of: $N_i = \underline{N} + a(GDP\ per\ capita_i)^{-b}$. Data from WDI.

Table D.4: Fertility and Real GDP per capita in 2000.

VARIABLES	(1) Fertility	(2) Log(Fertility)
\underline{N}	1.214*** (0.405)	
a	68.88** (32.37)	
b	0.454*** (0.0890)	
Log(GDP per capita)		-0.247*** (0.0160)
Constant		3.078*** (0.135)
Observations	181	181
R-squared	0.595	0.499

Notes: Data source: WDI for a cross-section of countries in 2000. The first-column estimates the following equation using non-linear least square (NLS): $N_i = \underline{N} + a(GDP\ per\ capita_i)^{-b}$. The second column (constant-elasticity case) estimates the following regression using OLS: $\log(N_i) = cste - b\log(GDP\ per\ capita_i)$. Standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1.

Results are shown in Table D.4. In our baseline specification, the asymptotic fertility rate \underline{N} is found strictly positive but below the replacement rate — even though one cannot reject $\underline{N} = 2$ at 5%. In the constant elasticity case, the OLS estimate of the elasticity b is 0.25 — in the range of estimates using micro data on individuals within a country (see Jones and Tertilt (2008) on US data).⁷⁸

Natural fertility rate in China. Based on the previous estimation, we assume that the natural

⁷⁸The estimate of fertility to income is remarkably robust for different years of cross-sections of countries. Using data on Chinese households not affected by the policy in the Census 1982, we also find a similar elasticity across Chinese households.

fertility rate in the model $N_t = 2n_t$ obeys the following fertility-income relationship:

$$N_t = \underline{N} + aw_{\gamma_{n,t}}^{-b} \quad (28)$$

with $w_{\gamma_{n,t}}$ the parental wage at date t , \underline{N} , a and b positive parameters corresponding to the estimation of Eq. 26. When differentiated over time, Eq. 28 can be rewritten:

$$N_{t+1} - \underline{N} = (N_t - \underline{N}) \left(\frac{w_{\gamma_{n,t+1}}}{w_{\gamma_{n,t}}} \right)^{-b} \quad (29)$$

Given our baseline estimate, we target $\underline{N} = 1.2$ and $b = 0.45$. In other words, the preference for fertility v_t starts at its steady-state value of 0.58 pre-1970 for which $N_{t < 1970} = 2.92$ and is set in the later periods to values such that Eq. 29 holds in equilibrium — all other parameters of the model are set to their baseline values of Table D.4.

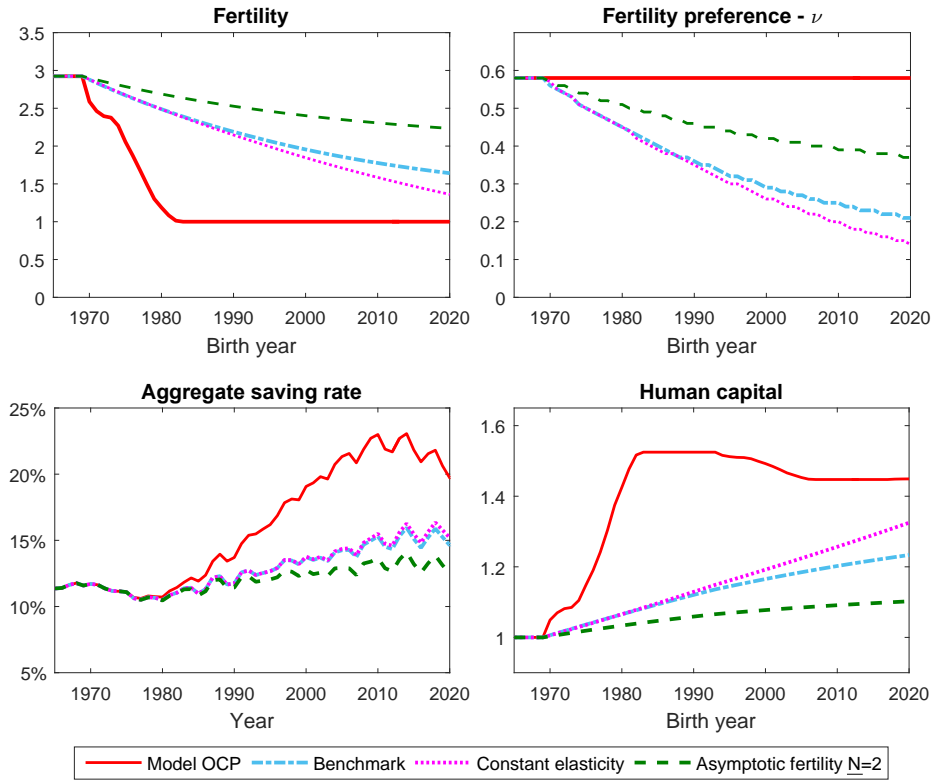
Given some uncertainty on the estimate of the asymptotic fertility rate, we provide a simulation where \underline{N} is set to 2, with b set to the same value — scenario corresponding to a constant population in the long-run. We also provide a simulation corresponding to the constant elasticity case with \underline{N} set to 0 and $b = 0.24$ (column 2 in Table D.4). The fertility rate under these alternative scenarios is shown together with the corresponding path of preferences v_t in Figure D.11 (upper-panel). Importantly to our results, the fall in fertility implied by Eq. 29 is much slower than under the one-child policy — in our baseline scenario, the natural fertility reaches 2 in the early 2000s.

Simulation of the quantitative model and results. Under such a path for v_t and the corresponding natural fertility rate, our model is simulated since 1970 — keeping all parameters but v_t to their baseline values. Figure D.11 (bottom-panel) shows the time-series of the aggregate saving rate and human capital H_t over the period 1970-2020 in the simulated model with endogenous fertility. Outcomes under the one-child policy are also shown for comparison purposes. In our baseline scenario, the aggregate saving rate increases by 5% over the period 1982-2014 compared to 11.6% under the one-child policy. A generation born in 2000 has a 17% higher human capital than a generation born before 1970 while the difference is about 50% under the one-child policy. In our less (resp. more) conservative scenario where \underline{N} is equal to 2 (resp. 0), the saving rate increases by 3.1% (resp. 5%) since 1982; the human capital of a generation born in 2000 is 8% (resp. 19%) higher than a generation born before 1970. When comparing the path with and without policy, we find that the policy contributes to about 35% (resp. 45%) of the overall increase in aggregate saving in the model over the period 1982-2014.

‘Two-children’ policy. For comparison purposes, we show the predictions of the model under a ‘two-children policy’. We implement a ‘two-children policy’ in the model by assuming that fertility declines linearly from 2.92 to 2 children per household over the period 1970-1978 and cannot exceed two children per household post-1978 — $n_{\max,t} = 1$ for $t \geq 1978$.⁷⁹ All other parameters of the model are set to their baseline value of Table 3. Figure D.12 shows the time-series of the aggregate saving rate and human capital H_t over the period 1970-2020 in the simulated model with endogenous fertility. Outcomes under the one-child policy are also shown for comparison purposes.


⁷⁹In our ‘two-children policy’ experiment, we assume that the fertility constraint is always binding (fertility equal to 2 starting 1978) while the natural fertility rate might have fallen slightly below 2 in the most recent period as shown in the previous experiments. However, as the natural fertility rate stays above 2 for most of the period in our experiments and close to 2 in the 2000s, results are almost identical if the constraint is not binding in the 2000s due to changes in fertility preferences. Results are also very similar quantitatively under an alternative assumption regarding the progressive decline over the period 1970-1978.

Figure D.11: Fertility, saving and Human Capital: natural fertility rate counterfactuals.



Notes: The left upper-panel of the figure plots the natural fertility rate predicted in our counterfactuals. The ‘benchmark’ line corresponds to our baseline (Eq. 29 with $\underline{N} = 1.21$ and $b = 0.45$). The ‘asymptotic fertility $\underline{N} = 2$ ’ line assumes fertility to be equal to the replacement rate in the long-run ($\underline{N} = 2$); the ‘constant elasticity’ line corresponds to the constant-elasticity case ($\underline{N} = 0$ and $b = 0.24$). The right upper-panel plots the corresponding path of ν_t in each of these scenarios. Given this path of ν_t and all other parameters set to their baseline values of Table 3, the bottom-panel shows the model’s predictions for the aggregate saving rate and the level of human capital H_t in the different scenarios for the natural fertility rate. The baseline simulations under the one-child policy (Model OCP) are shown for comparison purposes.

Figure D.12: saving and Human Capital: ‘Two-children’ policy counterfactual.



Notes: The figure the model’s predictions for the aggregate saving rate and the level of human capital H_t under a ‘Two-children’ policy (Model 2CP). The baseline simulations under the one-child policy (Model OCP) are shown for comparison purposes. In the ‘Two-children’ policy counterfactual, fertility declines linearly from 2.92 to 2 children between 1970 and 1978 and remains at 2 children thereafter ($n_{\max,t} = 1$ for $t \geq 1978$). All other parameters set to their baseline values of Table 3.

E Online Appendix: Data Treatment

E.1 Individual consumption estimation

The estimation procedure for age-saving profiles in China is explained in details in the online Appendix of Coeurdacier, Guibaud and Jin (2015). Here, we briefly describe the main methodology employed to disaggregate household consumption into individual consumption, and thereby estimate individual saving by age.

Projection Method. Our approach applies a projection method proposed by Chesher (1997, 1998) and Deaton and Paxson (2000). Essentially, the idea is to recover the consumption of each individual member of the household using cross-sectional variations in the composition of households as a source of identification. In practice, this is done by projecting household consumption on the number of household members belonging to various age groups, controlling for observable household characteristics. Following Chesher (1997), we conduct a non-linear least squares estimation of the following model for each year:

$$C_h = \exp(\boldsymbol{\gamma} \cdot \mathbf{Z}_h) \left(\sum_{j \geq 21} c_j N_{h,j} \right) + \epsilon_h,$$

where C_h is the aggregate consumption of household h , $N_{h,j}$ is the number of members of age j in household h , and \mathbf{Z}_h denotes a set of household-specific controls (income group, number of adults, number of children, uni- vs. multi-generational, etc.).⁸⁰ The estimated consumption of an individual of age j living in a household with characteristics \mathbf{Z}_h is then equal to $\exp(\hat{\boldsymbol{\gamma}} \cdot \mathbf{Z}_h) \hat{c}_j$. Details of the methodology, as well as robustness checks, are given in the online Appendix C.2. of Coeurdacier et al. (2015).

E.2 Empirical Counterfactual

Our empirical counterfactual aims at identifying the different channels through which having two children rather than one affect household saving. Four different effects comprising the macro-economic and micro-economic channels include (i) composition of population; (ii) composition of income and education; (iii) expenditure channel; (iv) transfer channel. We decompose the quantitative contribution of each of these different channels in Table 10, noting however that (iii) and (iv) are difficult to disentangle empirically.

Macro-channels.

Composition of Population. First, one needs to account for the shifts in the demographic composition. This involves multiplying the number of observations of individuals born after 1980 by a factor of 2 in the 2009 sample.⁸¹ Holding constant the age-saving profile, aggregate saving is now about 1.62 % lower under a ‘two-children policy’ due to the demographic composition effect.

Composition of Education and Income. Second, the incremental individual human capital that is at-

⁸⁰This assumes that individual consumption can be written as multiplicatively separable functions of individual age and household characteristics. The identification therefore relies on the restriction that the effect of household characteristics on individual consumption is independent of age.

⁸¹We correct the 2009 age distribution (UHS) to match the urban age distribution in 2009 as implied by the 2000 census. This correction has only a small effect on the estimates.

tributed to the one-child policy alters household saving to the extent that it alters the composition of income across age groups. As young individuals tend to save less, increasing their share of aggregate income relative to middle age individuals reduces the aggregate household saving rate. Therefore, we need to ‘purge’ the additional human capital caused by the policy. Using estimates of the twin-effect on education attainment provided in Table 8, we give young cohorts a 40 percent lower likelihood of attaining higher education under the two-children scenario. The impact on aggregate saving, holding everything else constant, is however very small—less than 0.4%. The effect being small is not surprising since it concerns only a small fraction of households in the whole sample at present. Note that the impact of higher education on aggregate saving due to the policy should switch sign in future years: the policy will increase the income share of middle aged individuals which have a higher propensity to save.

Micro-channels.

Expenditure and Transfers. Third, the imputed increase in expenditures associated with having an additional child is used to quantify the expenditure channel effect. Taking first education expenditures, all households with a child respectively under (resp. above) 15 years of age in the sample are given a 3.49% (resp. 6.71%) higher expenditure in education (as a share of household income), relying on the estimates from Table E.1 (Column 1). The overall effect of higher education expenditures leads to an additional 2.52% fall in the aggregate saving rate.

One can proceed by the same methodology to calculate the additional non-education related expenditures, remarking though that these effects kick in mostly during later stages of adulthood (Table E.1, column 2). We impute to all parents with financially dependent children (i.e. below 21 or below 25 and still students) a 1.58% higher non-education expenditure when under 45, and a 5.64% higher expenditures when above 45.⁸²

Taken all together, the incremental education and non-education expenditures lead to an additional 3.95% (= 2.52% + 1.43%) drop in the aggregate saving rate (see Table 10). Note that apart from education expenditures that are clearly devoted to children, the change in other expenditures when moving from one to two children is partly driven by the ‘expenditure channel’ and partly by the ‘transfer channel’. One cannot fully disentangle the two using this methodology, but we nevertheless believe that the impact on older parents’ of ‘other expenditures’ is likely to operate through the transfer channel.

A caveat is that older parents (in their late 40s and 50s) that were subject to the policy should also be affected by the ‘transfer channel’, even though their only child has left the household. This effect cannot be properly measured in the data since one can no longer observe whether parents had an only child or twins once the children have departed from the household (except in the 1 or 2 following years using the panel dimension). But if the twin effect on ‘non-education expenditures’ for parents above 45 is used as a proxy for their increase in overall expenditures, (treated) households in their late 40s to 50s (before retirement) with two children should incur an additional 6% (of household income) higher expenditure (in line with column (2) of Table E.1). This lowers the aggregate saving rate by an additional 1.57%. This channel is, however, less precisely estimated from the data and warrants a sensitivity analysis using more conservative estimates: assuming instead that additional expenditures

⁸²In Table E.1, we do not restrict the sample to nuclear households as in Tables 5 and 7. Estimates from Table E.1 are thus slightly different from the ones from Table 5 (Column 4) and Table 7 (Column 3). Our counterfactual applies to all households — the reason we use estimates from this larger sample. Education expenditures in Column 1 are also net of transfers for a child education in another city in order to ensure consistency with the model calibration.

Table E.1: Education and non-education expenditures for different age groups: Twin identification.

VARIABLES (in % of household income)	(1) Education exp.	(2) Non-education exp.
Twins < 15	0.0349*** (0.00597)	
Twins ≥ 15	0.0671*** (0.0154)	
Twins with parents < 45		0.0158 (0.0126)
Twins with parents ≥ 45		0.0564** (0.0243)
Controls	YES	YES
Observations	36,831	36,544
R-squared	0.123	0.172
Years Dummies	YES	YES
Prefecture Dummies	YES	YES

Notes: Data source: UHS (2002-2006). Restricted sample of households with an oldest child of age up to 21 years old. Outliers with saving rate over (below) 85% (-85%) of income are excluded. Controls include average age of parents, mother’s age at first birth, oldest child’s age, household income and a dummy for the multigenerational structure of the family. In column (1) a dummy for an oldest child above the age of 15 and in column (2) a dummy for parents above the age of 45. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

are 3% higher (rather than 6%) for older parents (between 45 and 60y without children below 21 or below 25 but still studying in the household), the aggregate saving rate falls by an additional 0.79%.

The combined effect of these channels summarized in Table 10 indicates that the aggregate saving rate would have been between 6.0% to 6.8% lower if China had implemented a (binding) ‘two-children’ policy—or, alternatively if the natural rate of fertility post 1980 had simply been two children per household. These estimates impute about a third of the increase in the aggregate saving rate in China to the one-child policy.

Discussion of the empirical counterfactual.

Twins versus two children. One may take caution with the empirical counterfactual as it assumes that having twins is similar to having two children sequentially. In what ways are twins different from two singleton children? One is that their arrival together may have been unanticipated, and the second is that there may be a difference in the degree of scale economies when having twins compared to having two children sequentially. With regard to the first issue, we mainly focus on older parents, so that the unanticipated dimension to the arrival of twins is less relevant—particularly since most of the expenditures relates to education at a later stage. Regarding economies of scale, there is evidence that they are larger for short birth spacing, particularly in childcare, and when children are of the same gender (see Newman (1983), Browning (1992), Rosenzweig and Wolpin (2000) for references; see also Rosenzweig and Zhang (2009)). For instance, apparel and room-sharing are found to be more likely when siblings are of the same gender. Thus, if anything, we tend to underestimate the expenditure channel when focusing on twins even though scale economies should not significantly differ between twins and two children households for most part of education costs.