THE ONE-CHILD POLICY AND HOUSEHOLD SAVING

Taha Choukhmane
MIT Sloan School of Management, USA

Nicolas Coeurdacier
SciencesPo, France

Keyu Jin
London School of Economics, UK

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. (JEL: E21, D10, D91)

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 three children per family in the late 1960s to just about one in the early 1980s. The

The editor in charge of this paper was Nicola Pavoni.

Acknowledgments: 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, and UC Berkeley for helpful comments. Nicolas Coeurdacier thanks the European Research Council (ERC) for financial support (ERC Starting Grant INFINHET 336748) and the SciencesPo-LSE Mobility Scheme. Choukhmane is affiliated at NBER. Coeurdacier is affiliated at CEPR.

E-mail: tahac@mit.edu (Choukhmane); nicolas.coeurdacier@sciencespo.fr (Coeurdacier); k.jin@lse.ac.uk (Jin)

Journal of the European Economic Association 2023 21(3):987–1032
https://doi.org/10.1093/jeea/jvad001

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 30 years (Li et al. 2017), with the average years of schooling increasing from 6.3 to 8.7 years for an adult aged 25 (Barro and Lee 2010).

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 three children were suddenly constrained to one. The reduction in expected transfers means that parents now have to save more on their own. Parents shift their investment in the form of children toward 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%–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 intergenerational 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.

Measuring the effect of the one-child policy on saving is challenging since cohort-level savings at different ages may have increased in China over the last decades for many reasons unrelated to the policy. Thus, instead of comparing in a reduced-form approach the saving behavior of cohorts differently affected by the policy, we provide a model-based quantification of the joint impact of the policy on human capital and saving. We validate the quantitative predictions using the births of twins as exogenous deviations from the one-child policy. More specifically, 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 based on a “twin experiment”.

Our theoretical framework enriches the standard life-cycle theory of saving with two additional elements: 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 how much to save for retirement. Children are costly but, at the same time, present 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 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.

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.

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 5–8 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 macro-economic implications. Storesletten and Zilibotti (2014) 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), 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.[3] 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 an identification strategy (Rosenzweig and Wolpin 1980).[4] 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 rather than through a social security system—the existing system leaving the majority of Chinese workers uncovered. Our baseline model treats these transfers toward the elderly as a social norm and thus exogenously given, contrary to Imrohoroglu and Zhao (2018). While their framework is richer in modeling transfers toward elderly to insure long-term care risks, ours emphasizes novel interactions

1. Modigliani and Cao (2004), Horioka and Wan (2007), and 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 (2018) goes further in inspecting the transfer channel through which fertility affects saving. They also provide comforting micro-evidence.

2. Some compelling explanations of the saving puzzle include: (1) precautionary saving (Blanchard and Giavazzi 2005, Chamon and Prasad 2010, and Wen 2011); (2) changes in income profiles (Song and Yang 2010); (3) gender imbalances and competition in the marriage market (Wei and Zhang 2011 and Du and Wei 2013); (4) demographics (Modigliani and Cao 2004, Horioka and Wan 2007, Curtis, Lugauer, and Mark 2015, Banerjee et al. 2014, and Imrohoroglu and Zhao 2018); (5) income growth and credit constraints (Coeurdacier, Guibaud, and Jin 2015), interacted with housing costs (Wang and Wen 2012, Bussiere et al. 2013, Wan 2015, and Lan 2019); and (6) reallocation of resources toward private firms (Song, Kjetil, and Zilibotti 2011). 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.

3. See also Boldrin and Jones (2002), Chakrabarti (1999), Cisno and Rosati (1996), Manuelli and Seshadri (2007), and Raut and Srinivasan (1994).

4. See Angrist, Victor, and Schlosser (2010) for references. Rosenzweig and Zhang (2009) use the birth of Chinese twins to measure the "quantity–quality" trade-off and find supporting evidence (see also Hongbin et al. 2008; Oliveira 2016; Qian 2013).

between fertility, human capital formation, and saving in the presence of old-age support. Our model also treats interest rates as exogenous, abstracting 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.[5]

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 are conducted in Section 5. Section 6 concludes.

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 Online 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.[6] 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 accounts for 88% of the increase between 1982 and 2014.[7]

5. Despite capital controls, China is also a semi-open economy where household saving is largely channeled abroad. A general equilibrium analysis may be found in Banerjee et al. (2014) and our related work (Coeurdacier et al. 2014).

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

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

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, Scharping 2003). The timing and the extent of enforcement of these policies varied across regions and significant discretion was given to local governments to implement them. In the late 1970s, the Chinese government shifted to a stricter approach, 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 3 years, the share of first-births in total births jumped from a 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 10 years later (Figure 1, bottom-panel). Crucially, 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 an 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 Online Appendix B that the one-child policy can account for the gradual decrease in fertility for parents who had children in the 1970s.[8] Additional evidence of the major role played by the policy in constraining fertility is provided in the same Online 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. 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

8. Assuming 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 Online 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.

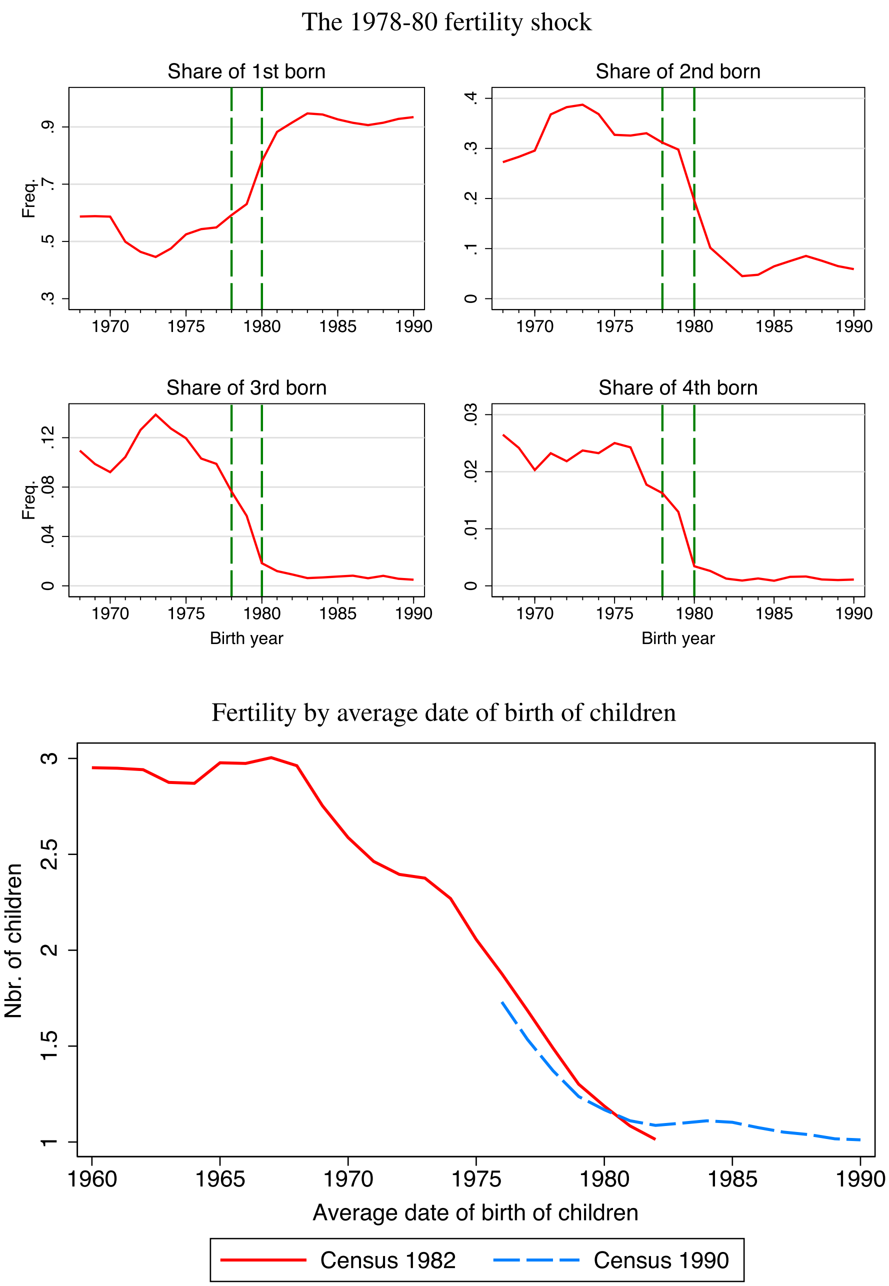

FIGURE 1. The one-child policy and fertility in urban China. The upper panel shows the number of births of a *n*th child divided by the total number of births in a given year. The vertical lines correspond to a two-children limit in 1978 and the one-child policy in 1980. The bottom panel shows the completed 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 Online Appendix A.

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).

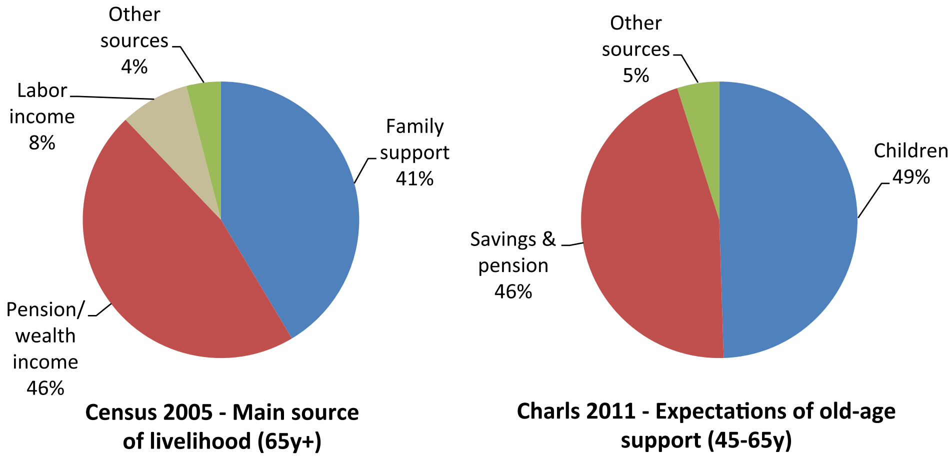

FIGURE 2. Main source of livelihood for the elderly (65+) in urban areas. Left panel, Census (2005). Right panel, China Health and Retirement Longitudinal Study (CHARLS 2011), urban households, whole sample of adults between 45and 65 (answer to the question: Whom do you think you can rely on for old-age support?).

older population (above 60) increases sharply only after 2015—when the one-child generation ages. In other words, the one-child policy leads first to a sharp fall in the share of young relative to middle-aged individuals, followed by a rapid rise in the share of the elderly only one generation later.

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 the 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 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).

China Health and Retirement Longitudinal Studyprovides 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. Old age support takes broadly two forms: financial transfers ("direct" transfers) and "indirect" transfers in the form of co-residence or other in-kind benefits. A total of 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 elderly's income (and up to 61% if one focuses on transfer receivers) (see Online Appendix Table A.1 for additional details). In addition, transfers (as a % of total income) are increasing in the number of children from 16% of parent's total income for households with one or two adult children to 80% for those with five or more children . The flip side of the story is that restrictions in fertility will therefore likely reduce the amount of transfers conferred to the elderly (see Online Appendix Table A.1). 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 toward their children according to the Chinese Household Income Project (CHIP) in 2002. 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.[9] 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.

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

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

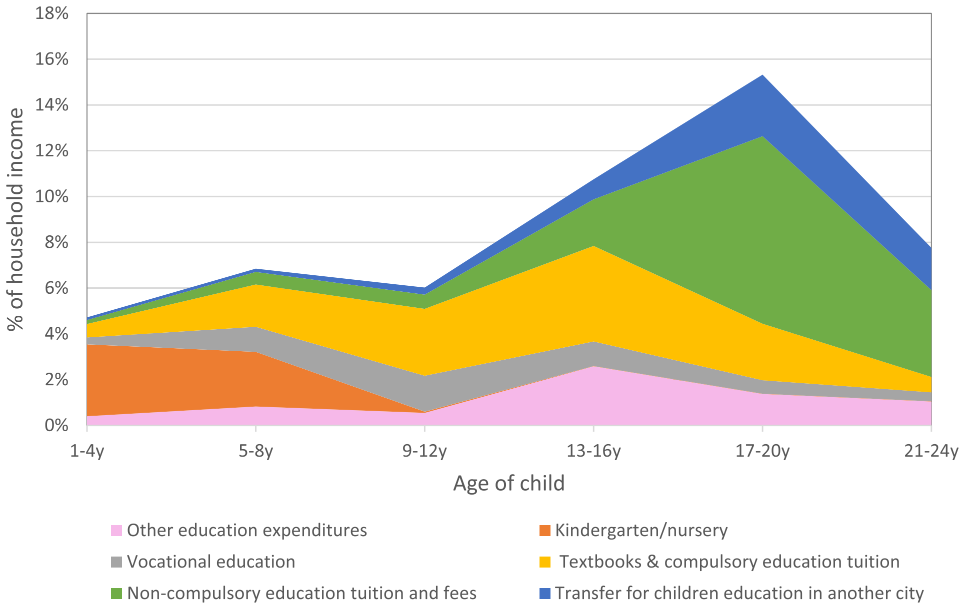

FIGURE 3. Education expenditures for a child, by age of the child (% of household income). Data source: Chinese Household Income Project (CHIP 2002). Sample restricted to urban households with an only child. This graph plots the average expenditure (as a share of household income) across education categories by the age of the only child.

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 with the notion that education investment is the main form of transfers toward children. After this age, children confer increasing amounts of transfers toward their parents—received by parents upon retirement (right panel). Considering co-residence (right axis) as an alternative form of transfers, children leave the parental household upon reaching adulthood (left panel).[10] For parents in their 60s–early 70s, the degree of co-residence falls less with parental age, remaining around 40%–50% as children are less likely to leave their parents at older age (right panel).

3. Theoretical Analysis

We develop an overlapping generations model with intergenerational transfers, endogenous fertility and human capital accumulation. The parsimonious model yields a tractable solution that serves two main purposes. First, it reveals the fundamental channels driving the fertility–human capital–saving relationships. Second, the model

10. Co-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.

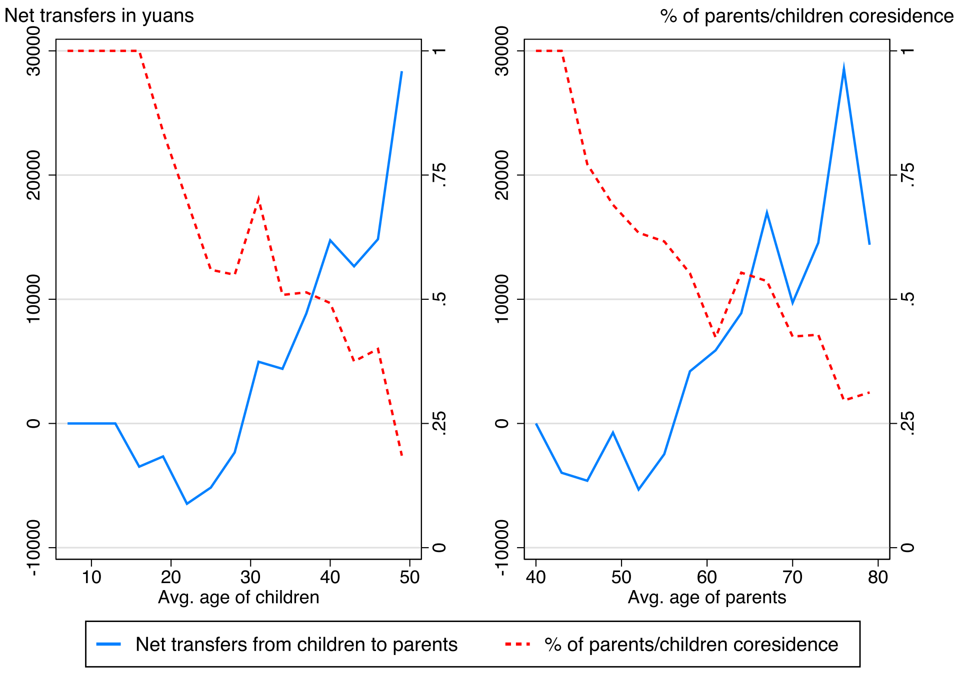

FIGURE 4. Timing of intergenerational transfers. China Health and Retirement Longitudinal Study (2008), whole sample of urban households. The left panel plots the average amount of net transfers of children to his/her parents (left axis) and the % of coresidence (right axis) by the average age of child. 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.

motivates our empirical strategy, showing how one can identify the impact of the one-child policy on human capital and saving through a cross-sectional comparison between two-children (twin) and only-child households. A quantitative version of the model is developed subsequently, although the main mechanisms are elucidated in the simpler model.

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. 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:

$$w_{y,t} = ez_t h_{t-1}^{\alpha} \quad \text{and} \quad w_{m,t+1} = z_{t+1}h_{t-1}^{\alpha}. \tag{2}$$

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, $n_t\varphi(h_t)w_{m,t+1}$, where $\varphi(h) = \varphi_0 + \varphi_h h$, $\varphi_0 > 0$ and $\varphi_h > 0$. 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 $\varphi_h h_t$ 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—capturing some crowding-out of individual transfers when more siblings are providing old-age support.[11] We treat these transfers as an institutional norm in China; children supporting their parents are not only socially expected, but

11. This crowding-out captures lower individual incentives to transfer when the amount transferred to the parents increases, or alternatively, some free-riding among siblings sharing the burden of transfers. It could also be related to a change of the social norm when the family size shrinks.

are even stipulated by law. The assumed functional form for transfers is analytically convenient, but its main properties are tightly linked to the data (see Section 4.2).[12]

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

3.2. Household Decisions and Model Dynamics

Consumption Decisions. Optimal consumption can be solved given fertility and human capital decisions.

ASSUMPTION 1. The young are subject to a credit constraint, binding in all periods,

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

Assumption 1 specifies that the young can borrow up to a constant fraction θ of the present value of future wage income.[13] Assumption 1 and the absence of bequests mean that the only individuals that optimize their saving are the middle-aged. From the individual optimization problem and constraints in equation (1), the asset holding of a middle-aged individual is

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

Equation (3) illuminates the link between fertility and saving: Parents with more children accumulate less wealth because they have less available resources for saving (term $n_t\varphi(h_t)$) and because they expect 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:

$$\frac{v}{n_t} = \frac{\beta}{c_{m,t+1}}\left(\varphi(h_t) - \mu_{t+1}\psi n_t^{\omega-1}\left(\frac{h_t}{h_{t-1}}\right)^{\alpha}\right)w_{m,t+1}, \quad (4)$$

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: It is equal to 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

12. In 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).

13. 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).

incentives to procreate depend on the factor μ_{t+1} defined as productivity growth relative to the gross interest rate. Higher productivity growth raises the number of children—by raising future benefits, while a higher interest rate increases the opportunity cost of child rearing.

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

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

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 equation (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}:

$$h_t = \left[\frac{\psi}{\omega\varphi_h}\frac{\alpha\mu_{t+1}}{h_{t-1}^{\alpha} n_t^{1-\omega}}\right]^{\frac{1}{1-\alpha}}. \tag{5}$$

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 generosity ψ of children toward parents also increases parental investment in them.

The optimal number of children n_t, combining equations (3), (4), and (5), satisfies

$$n_t = \left(\frac{v}{\beta(1+\beta)+v}\right)\left(\frac{1-\theta-\psi\frac{n_{t-1}^{\omega-1}}{\omega}}{\varphi_0+\varphi_h(1-\lambda)h_t}\right) \text{ with } \lambda = \frac{v+\omega\beta(1+\beta)}{\alpha v+\alpha\beta(1+\beta)}. \tag{6}$$

Equations (5) and (6) are two equations that describe the evolution of the two state variables of the economy $\{n_t; h_t\}$. Equation (5) describes the human capital response to a change in fertility n_t—with h_t decreasing in n_t. Equation (6) 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 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 $n_t = n_{ss}$.

Equations (5) and (6) are in steady-state,

$$\frac{n_{ss}}{1-\theta-\psi n_{ss}^{\omega-1}/\omega} = \left(\frac{v}{\beta(1+\beta)+v}\right)\left(\frac{1}{\varphi_0+\varphi_h(1-\lambda)h_{ss}}\right), \tag{NN}$$

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

Equation (NN) describes the higher demand for children (for $\lambda \geq 1$) when they have higher levels of human capital. Equation (QQ) illustrates the quantity/quality trade-off in children. Under the assumption that $\omega \geq \alpha$, this determines a unique stable steady-state of fertility and human capital toward which the economy converges (see Online Appendix C for a proof and a graphical illustration of the trade-off).

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, we must assume that $\alpha < 1/2$ for the model to converge as $T \to \infty$.[14]

3.3.1. Human Capital and Aggregate Saving.

Human Capital. The policy aimed at reducing the population also *increases* the level of per-capita human capital due to the quantity–quality trade-off (equation (QQ)). Therefore, the first generation of only child features a higher level of human capital than their parents (see Online Appendix C for a discussion and proof).

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 labor 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 Online Appendix C for details). Assuming constant productivity growth to interest

14. This assumption is needed 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 2003).

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 1. *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 Online 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

$$s_{t_0+1} - s_{t_0} = \underbrace{\frac{\left(n_{t_0-1} - n_{\max}\right) e}{1 + n_{\max} e} s_{t_0} + \frac{1}{1 + n_{\max} e} \theta\mu \left(n_{t_0-1} - n_{\max}\left(\frac{h_{t_0}}{h_{t_0-1}}\right)^{\alpha}\right)}_{\text{macro-channel (composition effects)}}$$
$$+ \frac{1}{1+n_{\max} e} \underbrace{\frac{\beta}{1+\beta}\left[\varphi_0\left(n_{t_0-1} - n_{\max}\right) + \left(\alpha + \frac{1}{\beta}\right)\frac{\psi\mu}{\omega}\left(n^{\omega}_{t_0-1} - n^{\omega}_{\max}\left(\frac{h_{t_0}}{h_{t_0-1}}\right)^{\alpha}\right)\right]}_{\text{micro-channel}}, \tag{7}$$

where the initial steady-state aggregate saving rate s_{t_0} is given in Online 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 (see Online Appendix C for a more detailled discussion). 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 $\varphi_0(n_{t_0-1} - n_{\max})$) 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 "$1/\beta$"). The increase in saving is mitigated by the human capital response. However, the increase in human capital is not large enough to offset the fall in fertility such that the middle-aged saving rate increases unambiguously with fewer children (see Online Appendix C for details).

3.3.2. Identification through "Twins". We next show theoretically how one can identify the micro-economic channel (over time) through a *cross-sectional* comparison between only-child households and twin-households. Proofs of these results are relegated to Online 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:

$$\frac{1}{2} \leq \left(\frac{h_{t_0}^{\text{Twin}}}{h_{t_0}} \right) = \left(\frac{1}{2} \right)^{\frac{1-\omega}{1-\alpha}} < 1. \tag{8}$$

The quantity–quality trade-off driving human capital accumulation can be identified by comparing twins and an only-child. 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 $\alpha \to \omega$).

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

$$s_{m,t_0+1} - s_{m,t_0+1}^{\text{Twin}} = \frac{\beta}{1+\beta} \left[n_{\max}\varphi_0 + \left(\alpha + \frac{1}{\beta} \right) \frac{\psi\mu}{\omega} n_{\max}^{\omega} \left(\frac{h_{t_0}}{h_{t_0-1}} \right)^{\alpha} \left(2^{\frac{\omega-\alpha}{1-\alpha}} - 1 \right) \right] > 0.$$

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 equation 7). 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.

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

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

Proof. See Online Appendix C. □

If the unconstrained fertility is two children per household, then we can identify the micro-economic impact of the policy—by comparing the saving rate of a middle-aged parent with an only child to the one of a parent with 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 in China prior to the policy change). That is, if $n_{t_0-1} > 2n_{\max}$, then

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

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.

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. However, 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 study the potential evolution of fertility in the absence of policies in the context of our quantitative model of Section 4.

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 (Allen, Qian, and Gu 2015; Song, Kjetil, and Zilibotti 2011; and Song et al. 2015). While the 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. 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 Overlapping Generations (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 provides directly testable evidence at the micro-level 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 $\gamma = \{1, 2, \ldots, \gamma_d\}$ coexist in the economy in each period. The timing of the events that take place over the life cycle 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 $\overline{\gamma}$, he is rearing and educating children while making transfers to his elderly parents. He reaches old age at age $\overline{\gamma}$, with $\gamma_n < \overline{\gamma} \leq \gamma_d$—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.[15] Our baseline abstracts from social security transfers and takes old age-support as given. Extensions of the baseline model in these dimensions are provided in Online Appendix D.3.[16]

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

$$U(t) = v \log(n_{t+\gamma_n}) + \sum_{\gamma=\underline{\gamma}}^{\gamma_d} \beta^{\gamma-\underline{\gamma}} \log(c_{\gamma,t+\gamma}), \tag{9}$$

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 $\gamma \in \{\gamma_n, \ldots, \gamma_n + \gamma_e\}$ are a fraction $\varphi_\gamma n_{t+\gamma_n}$

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

16. The baseline without social security is arguably not too far from the reality of the majority of Chinese urban households—due to the very low coverage rates of the existing social security system, as well as its falling generosity for covered workers over the period considered. Further details on Chinese social security are provided in Online Appendix D.3.2.

of the agent's wage income $w_{\gamma,t+\gamma}$. The discretionary education costs are borne at the same age and are a fraction $\varphi_{\gamma,h} h_{t+\gamma_n} n_{t+\gamma_n}$ 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 $\gamma \in \{\overline{\gamma}-\gamma_n, \ldots, \gamma_d - \gamma_n\}$ and are a fraction $\psi(n_t^{\omega-1}/\omega)$ of the wage income.[17] When old, at age $\gamma \geq \overline{\gamma}$, the agent receives transfers from his $n_{t+\gamma_n}$ children equal to $\psi(n_{t+\gamma_n}^{\omega}/\omega) w_{\gamma-\gamma_n,t+\gamma}$. 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:

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

where $\mathbf{1}_{\{x \leq \gamma \leq y\}}$ is equal to one if $\gamma \in \{x, \ldots, y\}$ 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:

$$a_{\gamma,t+\gamma} = w_{\gamma,t+\gamma} - c_{\gamma,t+\gamma} - T_{\gamma,t+\gamma} + R a_{\gamma-1,t-1+\gamma} \qquad \gamma \in \{\underline{\gamma}, \ldots, \gamma_d - 1\}, \tag{10}$$

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

$$a_{\gamma,t+\gamma} \geq -\theta \frac{w_{\gamma+1,t+\gamma+1}}{R} \quad \gamma \in \{\underline{\gamma}, \ldots, \gamma_d - 1\}. \tag{11}$$

Fertility Constraints. Fertility policies require that

$$n_t \leq n_{\max,t}, \tag{12}$$

where $n_{\max,t}$ captures fertility policies at every date t. If at date t, agents can freely choose fertility, then $n_{\max,t} \to \infty$. 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 $\{c_{\gamma,t+\gamma}\}_{\gamma \in \{\underline{\gamma},\ldots,\gamma_d\}}$, 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 $U(t)$ (equation 9), subject to a sequence of instantaneous budget constraints (equation 10), credit constraints (equation 11), and fertility constraints (equation 12). This characterizes consumption dynamics across age, as well as the dynamics of fertility and human

17. The baseline model assumes exogenous transfers driven by a social norm. Online Appendix D.3.3 provides an extension with endogenous transfers driven by a warm-glow motive.

TABLE 2. Calibration of model parameters.

Parameter	Main target (data source)	Value
$R - 1$ (annual)	Average real interest rate, 1979–2013 (details in Online Appendix D.2)	5.3%
g_z (annual)	Real wage growth (UHS)	6.1%
α	Mankiw, Romer, and Weil (1992)	0.37
υ	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 Figures 5 and 6
φ_γ	Compulsory education expenditures across ages in 2002 (CHIP)	and details in
$\varphi_{\gamma,h}$	Discretionary education expenditures across ages in 2002 (CHIP)	Online Appendix D.2

capital $\{n_t, H_t\}_{t>0}$ given initial conditions $\{n_0, H_0\}$. Details of the solution are provided in Online Appendix D.2.[18]

4.2. Data and Calibration

Timing. Agents live for 20 periods, where a period lasts 4 years. They start working in the sixth period (ages 21–24) and have children in the seventh (ages 25–28)—in line with the data.[19] They enter old age in period 16 (ages 61–64), age at which males retire in China. Figure D.1 in Online Appendix D.1 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 2 (details in Online Appendix D.2). Data used in the calibration are described in Online Appendix A.

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 for productivity growth g_z, as wage growth occurs

18. The model can be solved analytically if the credit constraints are not binding for ages $\gamma \geq \gamma_n$ (see Online Appendix D.2)—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.

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

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 Online Appendix D.2 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; Sianesi and van Reenen 2003).[20]

Age Income Profile. We calibrate the experience parameters $\{e_\gamma\}_{\gamma \geq \underline{\gamma}}$ to labor income by age group, provided by UHS data. The first available year for which individual labor 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 5.[21]

Real Interest Rate. In the spirit of Curtis et al. (2015) and 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 Online Appendix D.2 for details).

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 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, υ, to target $n_{ss} = n_{t<1970} = 2.92/2$. 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 Online 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. Post-1980, fertility is constrained by the one-child policy: $n_{\max,t} = 1/2$ for $t > 1980$.

20. Using equation (8), 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.

21. Calibrating experience parameters e_γ on the sole cross-section of 1992 data could mix age-effects and cohort-effects. Robustness checks discussed in Online Appendix D.2 show that it is not the case.

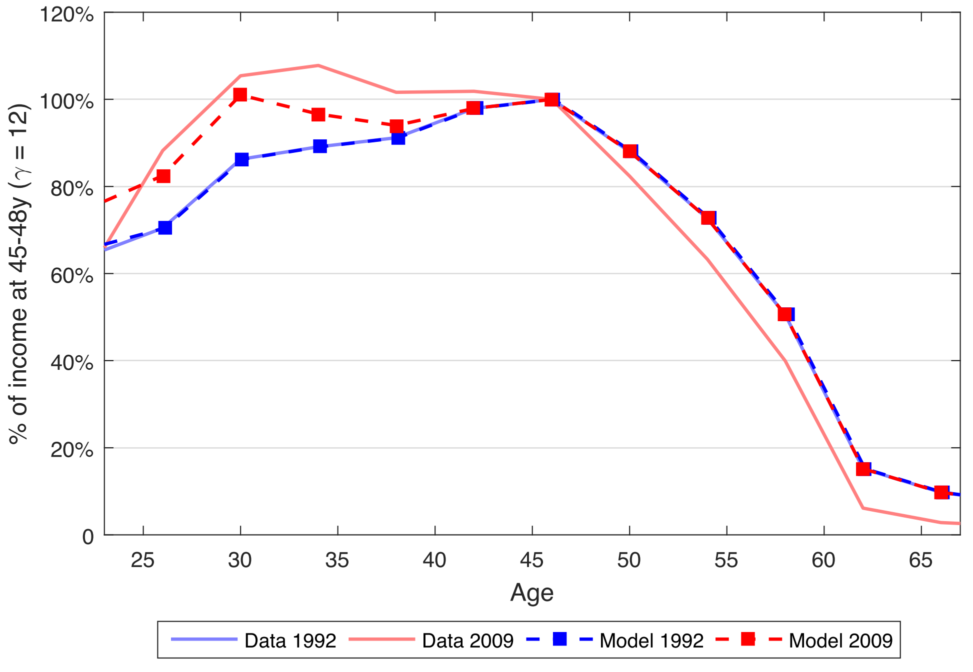

FIGURE 5. Age income profiles in 1992 and 2009. Model versus Data. This figure plots the model-implied labour income profiles by age in 1992 and 2009 and its data counterpart. Data source: UHS, 1992 and 2009. Wages includes wages plus self-business incomes. The profile in 1992 is used to calibrate experience parameters $\{e_\gamma\}_{\gamma \geq \underline{\gamma}}$. Parameter values for the model's simulations are provided in Online Appendix D.2.

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).[22] 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.[23]

Old Age Support. Two parameters govern transfers to parents, ψ and ω. The first captures the generosity toward parents in the economy; the latter captures the

22. Using the 1982 Census, we cannot reliably estimate the size of cohorts born before 1902 (i.e., aged above 61 in 1964). We therefore leave the age bins 61–64 to 77–80 undefined in 1964. This is unimportant, however, for our purposes because: (i) these agents do not make human capital decisions for the cohorts affected by the one-child policy; and (ii) we focus on aggregate saving starting 1982, at which point they are no longer alive.

23. Our model fits the distribution of population in the later years reasonably well (see Online Appendix D.2). However, it predicts age-groups of older individuals larger than in the data as it does not feature mortality before age γ_d.

crowding-out of individual transfers when the family size increases. We first estimate ω empirically.

Estimation of ω and Validation of the Transfer Function. China Health and Retirement Longitudinal Study 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 $\psi n^{\omega-1}/\omega w$. Details and results of the estimation are provided in Online Appendix D.2 (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 is linked to the overall level of transfers toward the elderly. Direct measurement of ψ based solely on measured transfers from CHARLS gives a low value for ψ, around 4%–5% for $\omega = 0.65$.[24] 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. Importantly, if one takes pecuniary transfers toward 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.

Computing Age-Saving Profiles. To set the remaining parameters, we target the saving rate by age in 1986. Age saving profiles are usually computed at the household level by age of the household head. As shown in Coeurdacier, Guibaud, and Jin (2015), such a measure might be inaccurate in presence of multigenerational households due to selection and aggregation biases. Thus, we follow their empirical strategy based on Chesher (1998) to estimate age saving profiles by age of individuals (see Online Appendix E.2 for details).

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

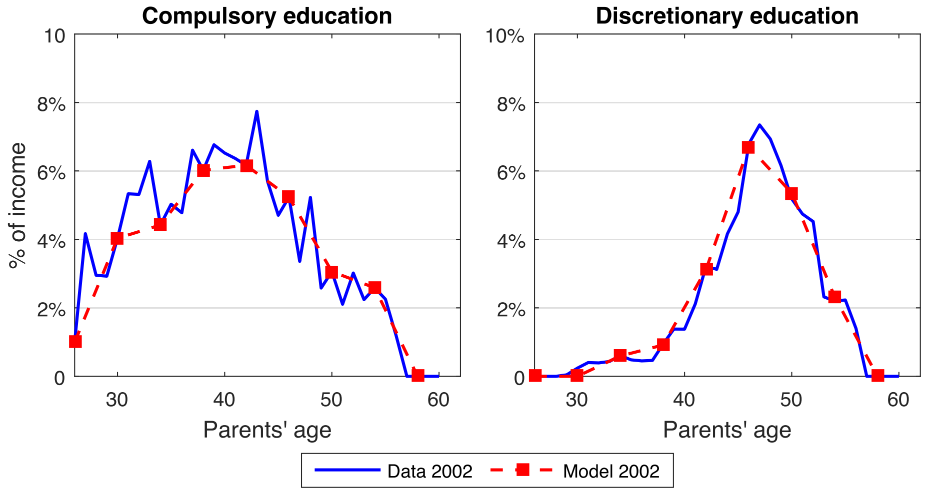

FIGURE 6. Education expenditures per child by age of parents in 2002. Model versus Data. This figure plots education expenditures by age of parents in 2002 in the data and in the model (in % of income). The left-panel shows compulsory education costs per child and the right panel shows discretionary education costs. Parameter values for the model's simulations are provided in Table 2 and detailed in Online Appendix D. The data counterpart is computed using CHIP 2002 (see Online Appendix A).

Parameters $\{\beta, \psi, \theta\}$ *and Education Parameters* $\{\rho; \varphi_\gamma; \varphi_{\gamma,h}\}_\gamma$. Our calibration strategy jointly determines the parameters $\{\beta, \psi, \theta\}$ and the education parameters $\{\rho; \varphi_\gamma; \varphi_{\gamma,h}\}_\gamma$ 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 $\varphi_{\gamma,h}$).[25] The fraction of wage income spent on compulsory education costs at a given age pins down the parameters $\{\varphi_\gamma\}_{\gamma\in\{\gamma_n,\ldots,\gamma_n+\gamma_e\}}$. As discretionary costs are very close to zero up to the age 10 of the child (Figure 6), we set $\varphi_{\gamma,h} = 0$ for $\gamma \leq 8$ (age 29–32).[26] 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 Online Appendix D.2). Based on this analytical expression, we show that for each value of the parameter ρ, there is a unique combination of the parameters $\{\varphi_{\gamma,h}\}_{\gamma\in\{\gamma_n,\ldots,\gamma_n+\gamma_e\}}$ such that the rate of change of discretionary costs between two ages matches its data

25. These 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.

26. Education costs are paid by parents until age 53–56 years and $\gamma_e = 7$.

counterpart in 2002. For a given ρ, the parameters $\{\varphi_{\gamma,h}\}_\gamma$ 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 $\{\varphi_\gamma; \varphi_{\gamma,h}\}_\gamma$, we search for the remaining parameters $\{\beta, \psi, \theta, \rho\}$ over a grid Ł 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 $\{\beta, \psi, \theta, \rho\} \in$ Ł to minimize the following distance:

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

where $s_{\gamma,1986}^m$ (respectively $s_{\gamma,1986}^d$) is the model predicted saving rate at age γ in 1986 (respectively the saving rate at age γ in the 1986 data); $educ_{\gamma,2002}^m$ (respectively $educ_{\gamma,2002}^d$) is the model predicted discretionary education spending as a share of wage at age γ in 2002 (respectively the discretionary education spending as a share of wage at age γ in the 2002 data); and λ_γ^s and λ_γ^{educ} are weights on different age groups summing to 1 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 7, upper panel).[27] 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 $\{\rho; \varphi_\gamma; \varphi_{\gamma,h}\}_\gamma$ fitting data on education spending in 2002 extremely well (Figure 6). The minimization leads to the following parameter values: $\beta = 0.99$ (annual basis); $\psi = 9\%$; $\theta = 0\%$; $\rho = 0.2$—the corresponding education costs $\{\varphi_\gamma; \varphi_{\gamma,h}\}_\gamma$ parameters being shown in Online Appendix D.2. The discount rate β is admittedly high though still in the ballpark of related papers.[28] Credit constraints are found to be very tight, in line with the low dissavings of young households and the low level of household debt.[29] Importantly,

27. As our sensitivity analysis shows (see Online Appendix D.3.1), 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.

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

29. The very low levels of household debt in China (about 10% of GDP in 2008) warrants a choice of a low θ to limit the ability of young households to borrow against future income. Our baseline calibration

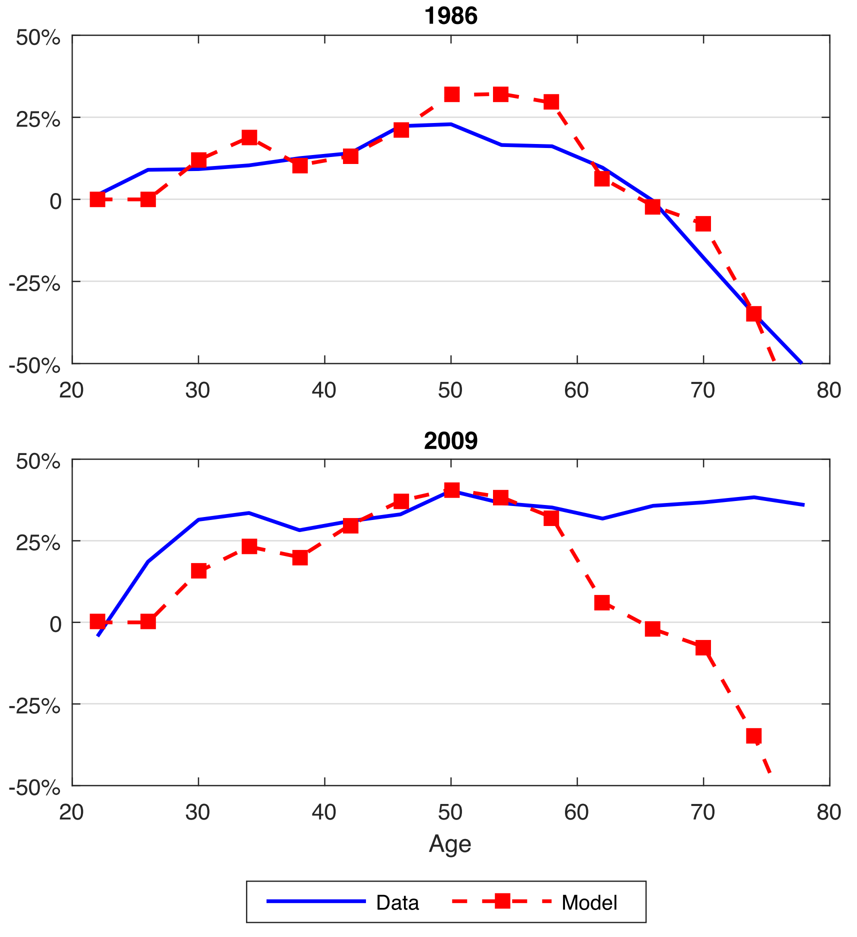

FIGURE 7. Age-saving profile in 1986 and 2009. Model versus Data. This figure plots age-saving profiles in 1986 and 2009 in the data and in the model. Parameter values for the model's simulations are provided in Table 2 and detailed in Online Appendix D. The data counterpart is estimated using UHS data (see Online Appendix E.2 for details on the estimation procedure).

the resulting value for the transfer parameter ψ is in line with Banerjee et al. (2014) and in line with data on pecuniary transfers toward parents living in another city.

Sensitivity and Extensions. Sensitivity with respect to the main parameters of the model is relegated to Online Appendix D.3.1. Online Appendix D.3.2 provides sensitivity to the presence of social security—the results in the following section remain largely unaffected under various scenarios regarding the system's generosity. Online

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 Online Appendix D.3.1 for sensitivity analysis.

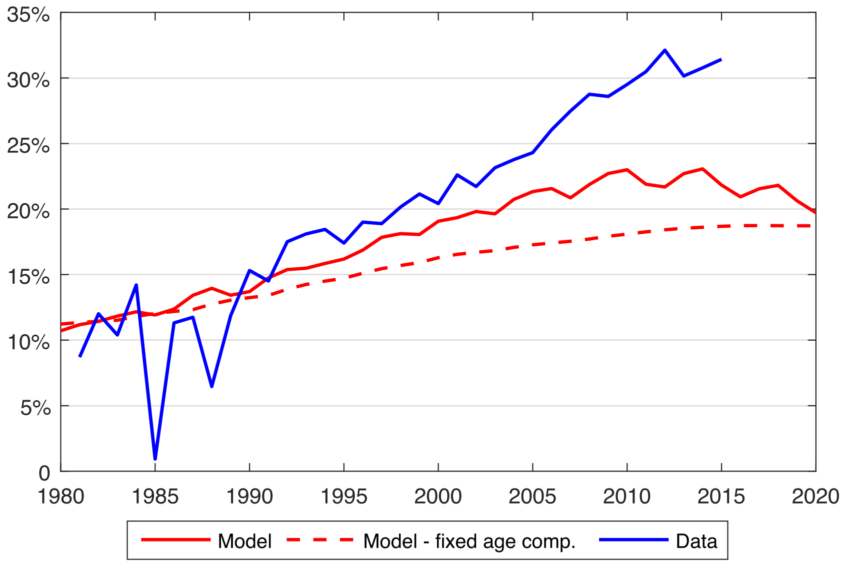

FIGURE 8. Aggregate household saving rate: Model versus Data. Data source: CEIC Data (using UHS). The model implied aggregate saving rate simulates the fertility policies using the calibration of Table 2.

Appendix D.3.3 develops an extension of the baseline model where transfers toward elderly parents are made endogenous through a warm-glow motive. Results are robust to this extension to the extent that transfers of siblings partly crowds out own individual transfers as exogenously captured by the transfer function in the baseline model.[30]

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. Household Saving.

Aggregate Saving. Figure 8 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 percentage points over the period 1982–2014, about 60% of the increase in the data. This is an *upper-bound* of

30. The extension with endogenous transfers generates this feature with two crucial ingredients: (i) at the margin, the warm-glow utility benefit from individual transfers toward parents decreases when siblings transfer more; and (ii) siblings do not coordinate their actions when deciding the amount transferred. Although endogenous transfers do not take the same functional form as in the baseline, the same properties holds: Transfers increase with (permanent) income and decrease with respect to the number of siblings (see Online Appendix D.3.3).

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 4.3.3 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. 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 8)—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 percentage points 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 percentage points rise over the period 1982–2014 in the baseline calibration ($\psi = 9\%$) compared to a 10 percentage points 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]

Saving by Age Groups. Beyond the trend in aggregate saving, we explore more micro predictions of our model for saving—comparing the saving rate of a given age-group implied by the model to its data counterpart.[32] Figure 7 compares age-saving profiles in 1986 (targeted) to 2009, in the data and in the model. Data shows an upward shift in the age-saving profile for all age groups but the youngest ones between 1986 and 2009. The increase in the saving rate for the middle-aged individuals (aged 30–50) lines up relatively well with the model's predictions, where it results from both a fall in expenditures on children and a fall in expected future receipts of transfers. Clearly, the model cannot account for the large increase in savings of the oldest age-groups as they were mostly unaffected by the policy. This increase for the elderly, and to some extent at the younger ages, constitutes the bulk of the increase in aggregate saving that the model cannot capture. While explaining such an increase at old-age is beyond the scope of the paper, rising longevity and rising health risks (together with a low coverage

31. See Online Appendix D.3.1 for sensitivity analysis with respect to ψ. Note that in order to match the level of aggregate saving 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 percentage points.

32. Alike for the 1986 cross-section, the average saving rate in an age-group at a given date is measured using UHS, correcting for the presence of multigenerational households (see Online Appendix E.2).

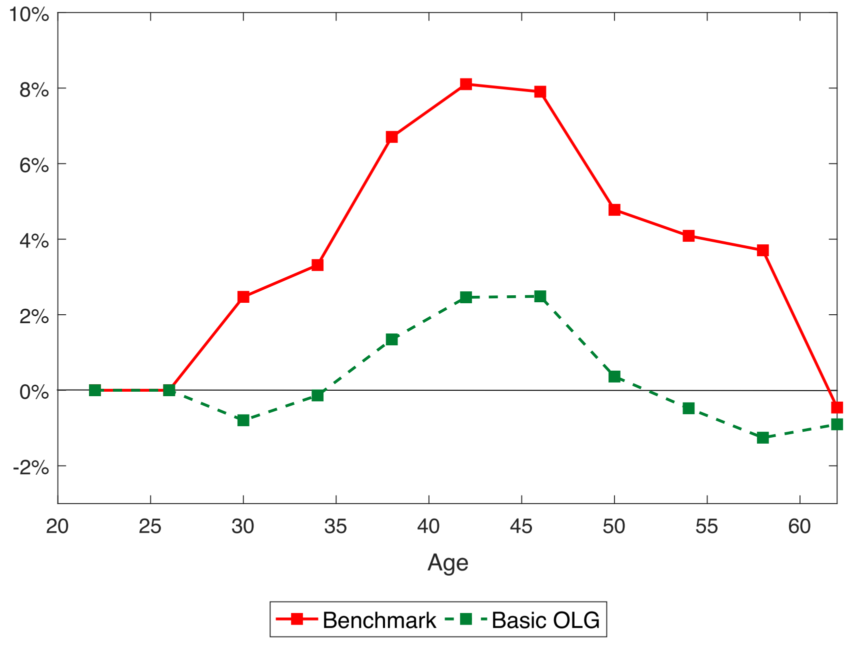

FIGURE 9. Difference in saving rates by age between parents of an only child and parents of twins. Model Predictions. This figure plots the model-implied difference in saving rates between parents of an only child and parents of twins in 2006 at different ages: $(s^{m}_{\gamma,2006} - s^{m,\text{Twin}}_{\gamma,2006})$. Two cases considered: our baseline calibration and a standard OLG model in which old age support and human capital accumulation are absent. Parameter values provided in Table 2.

of health insurance) are natural candidates (De Nardi et al. 2010). As other factors might have increased the savings of individuals at different ages independently of the one-child policy, we aim to isolate the role of fertility restrictions using cross-sectional comparisons of savings between parents of twins and parents of only child.

Saving in Only Child and Twins Households. A validation of the model's quantitative performance would rely on its ability to mimic differences in saving rates for parents of only child versus parents of twins. Figure 9 plots the predicted 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,[33] $\left(s^{m}_{\gamma,2006} - s^{m,\text{Twin}}_{\gamma,2006}\right)_{\gamma=\{\underline{\gamma},\ldots,\overline{\gamma}\}}$. 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.

To disentangle further the micro-channels, Figure 9 also displays the difference in saving rates between parents of an only child and parents of twins in a standard OLG model without old-age support. In this standard OLG model, only the *expenditure*

33. We use the prediction in 2006 as the data counterpart in our sample of twins covers the years 2002–2009. Results using other years over this period are very similar.

channel is operative.[34] The standard OLG model predicts much smaller differences in saving rates across all ages. The transfer channel thus appears quantitatively large in the model. Another important discrepancy between the two models concerns individuals in their 50s. Due to consumption smoothing, lower expenditures on children earlier in life release more resources for consumption when children no longer live in the household. Thus, the standard OLG model predicts lower saving rates for these age groups in households with fewer children, while our model predicts the opposite due to the *transfer channel*.[35]

These differences of saving rates between parents of an only child and parents of twins are at the heart of the empirical strategy developed in Section 5—investigating this difference in the data provides a clear test of the quantitative properties of our model.

4.3.2. Human Capital.

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 of an only child is 53% higher than the one of an individual born pre-policy in the late 1960s—translating into a wage increase of 17%. While the mapping between the model implied human capital and data is not straightforward, the number of years of schooling of the only-child generation born in the early 1980s is 1.5 years higher than a generation born in the late 1960s in urban China (see Online Appendix E.1 for details). Using a standard value of 10% of return to an additional year of schooling estimated in a Mincerian regression,[36] this translates into a wage increase of the only-child generation of 15%, fairly close to the model counterpart. Thus, once converted into wage increases, the model generates an increase in human capital close to its data counterpart. In line with these findings, the increase in human capital of the only-child generation explains a large fraction of the faster wage increase of young adults and the model generates *endogenously* a significant portion of the flattening of the age income profile observed in the data in 2009 (Figure 5).

Human Capital Of Only-Child versus Twins. 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 model-predicted effect of the policy if the natural fertility rate is around two. Differences in education spending

34. Education 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.

35. The transfer channel can be identified by investigating the saving behavior of parents after children have left the household. Banerjee, Xin, and Qian (2010), and Banerjee et al. (2014), using the partial implementation of fertility restrictions in the 1970s, compare the saving behavior of (treated) individuals in their 50s to (not-treated) individuals in their early 60s in 2008: The latter save on average about 10% less than the former. Our model implied difference (not shown) is very similar in magnitude.

36. Details of the Mincerian regressions using UHS data are provided in Online Appendix E.1. Standard values for the return of an additional schooling year hover between 6% and 13% (Card 1999, Psacharopoulos and Patrinos 2018).

and attainment between twins and only child are additional testable implications that motivate our subsequent empirical strategy.

4.3.3. 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 effect of the one child policy (as it assumes that the natural fertility would have stayed constant). Ideally, one would like to know how much these variables would have increased in the absence of any fertility policies. 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 under different hypotheses for the path of natural fertility. A first approach 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 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 Online Appendix D.4.

"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 of two children for $t \geq 1978$.[37] All other parameters of the model are set to their baseline value of Table 2. Under such a policy, the quantitative model predicts a 6.2 percentage points lower aggregate saving rate in 2014 than that under the one-child policy—about a third of the increase in the aggregate saving rate over the last 30 years. 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 conservative lower-bounds as fertility falls to 2 as early as 1978 in this simulation.

Natural Fertility Rate. With a constant preference for fertility v, the counterfactual fertility rate without constraints remains at its pre-1970 value—about three 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.[38] 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

37. We 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 Online Appendix B).

38. We assume that the link between fertility and income is driven by preferences v, which depends 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).

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:

$$N_t = \underline{N} + a w_{\gamma_n,t}^{-b}, \tag{13}$$

where the asymptotic fertility rate $\underline{N}$ and the parameters a and b are estimated in the cross-section of countries in 2000—details are provided in Online Appendix D.4. We then simulate our quantitative model assuming the path of v_t for which the fertility-income relationship of equation (13) holds—keeping all other parameters to their baseline value.[39] We find that the natural fertility rate falls 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 percentage points compared to more than 11 percentage points—implying that the one-child policy accounted for 35% of the observed saving rate increase.

Welfare Implications. Using our counterfactuals, we compute the welfare of different generations under the one-child policy or under a scenario where fertility is unconstrained. We do so under different scenarios for the natural fertility rate (status-quo to its initial value or downward trend due to rising income). Details of the results are relegated to Online Appendix D.5. Although quantitative results depend on the implied path of natural fertility, we find that fertility restrictions have redistributive welfare effects across generations across all simulations. The very first generations of parents subject to the one-child policy (born around 1960) are unambiguously hurt by the policy—their optimal level of children being constrained. However, for the later generations, the welfare effect of the policy is ambiguous. The first generations of only child (born around 1985), were also hurt as they could not freely choose their fertility. But they also benefited from the policy through a higher level of human capital investment of their parents. In our counterfactuals, we found that the latter effect dominate such that the generations of only child benefited from the policy (Table D.6 in Online Appendix D.5). Note that, once the policy ends, the very first generations able to choose freely their fertility are unambiguously better off due to their high human capital combined with unconstrained fertility decisions. These results show that fertility restrictions can be welfare improving in our framework, although it crucially depends on the welfare weights attributed by the planner to different generations as discussed in Online Appendix D.5. This is so because the level of human capital is inefficient in our framework. When parents decide the human capital of their children, they internalize their private benefits in the form of later transfers but do not take into account the welfare gains for their children.

39. We provide sensitivity analysis for the natural fertility rate around this baseline scenario: a scenario where the asymptotic fertility rate $\underline{N}$ is set to the replacement rate of two—above our estimated baseline but within the 5% confidence interval; a second scenario assumes a constant elasticity to income ($\overline{N} = 0$). See Online Appendix D.4 for details.

5. "Twin" Tests: Model versus Data

Section 3 showed how one can identify theoretically the micro-channel by comparing two-children (twin) households to only-child households. Using this 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 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).[40] 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. 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.

5.1. *Estimates of the "Twin Effect"*

Data used are described in details in Online Appendix A. A limitation is that one observes children only when (1) residing in a household, (2) when residing outside but remaining financially dependent, or (3) in the years just following their departure using the short panel dimension of the survey. This 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 older parents living alone.[41]

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

40. While 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). 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.

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

following regression is performed for a household h living in province p at a date $t = \{1986, 1992, ..., 2009\}$:

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

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-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; and $Z_{h,t}$ is a set of household level control variables—in particular, the (log of) age of parents and children. By including both age controls and year dummies, our regressions control for age effects and cohort effects. 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 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.

Columns (1)–(3) in Table 3 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 5–6 percentage points less than household with an only child. The magnitude is similar under different specifications and across samples.[42]

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 children after the implementation of the one-child-policy.[43] 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) and (6), we compute an alternative and more

42. In 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.

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

TABLE 3. Household saving rate: twin identification.

Variables	(1) Saving rate	(2) Saving rate	(3) Saving rate	(4) Saving rate	(5) Saving rate included Education transfers	(6)
Oldest child	Up to 18 years	Up to 18 years	Up to 21 years	Up to 18 years	Up to 18 years	Up to 21 years
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***	–0.0540***	–0.0566***	–0.0717***	–0.0839***	–0.0789***
	(0.0106)	(0.00971)	(0.00937)	(0.0123)	(0.0127)	(0.0122)
Twins born ≤ 1980	0.0177	0.0143	0.0197			
	(0.0173)	(0.0162)	(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	1,00,236	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 and 1992–2009). 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.Columns (5) and (6) include education transfers to children living in another city as part of consumption expenditures when computing household saving. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 4. Savings and expenditures for different age groups: twin identification.

Variables (in % of household income)	(1) Saving rate	(2) Saving rate	(3) Non-education expenditure	(4) Non-education expenditure
Twins	−0.0839***	−0.0655***	0.0360***	0.0195
	(0.0127)	(0.0137)	(0.0132)	(0.0144)
Twins with parents ≥ 45		−0.110***		0.0841**
		(0.0347)		(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) and (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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

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 visible at older age as shown in Section 4.3.1. 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 4 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 and 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. Saving differences between twins and only child: identification on "movers".

VARIABLES Oldest child	(1) Saving rate Up to 30 years birth ≥ 1980	(2) Saving rate Up to 30 years birth ≥ 1980
Adult twins left the household	−0.0920	−0.0910
	(0.0728)	(0.0728)
Adult singleton left the household	0.0698***	0.0708***
	(0.0117)	(0.0119)
Twins	−0.0498***	−0.0546***
	(0.00976)	(0.0109)
Twins 18–30 years		0.0189
		(0.0236)
Singleton 18–30 years		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). 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

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.[44] 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 UHS follows a given household for, at most, three consecutive years and "movers" constitute a small fraction of our sample of twins (about 20 observations).[45] Results are shown in Table 5 using the sample of households with children. Column (1) shows 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). Column (2) checks that our findings are not driven by the older age of "movers". 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

44. The "expenditure channel" generates *higher* saving rates of families with twins, once they have left (Figure 9).

45. Due to the lack of "movers" in the twins sample, we have to consider households in which one or two children have left.

different from zero). Most importantly, households with an only child still save more than twin households once a child has left.

Selection and "Artificial" Twins. Twins born after the one-child policy could potentially be "artificial" or "man-made" (Huang et al. 2016). If true, this is an issue if families with "artificial" twins have a different propensity to save/educate—after controlling for observable factors such as 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. We also investigated differences between only-child and twin parents across observable characteristics over time. We do not find that parents of twins are different in terms of education, income, or age at different periods—comforting our identification strategy.[46]

Quantity–Quality Trade-off. A quantity–quality trade-off is immediately visible from the evidence in Figure 10: The per-capita education expenditure on a twin is 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

$$\frac{\exp^{\text{Educ.}}_{h,p,t}}{n_{h,t}} = \alpha_t + \alpha_p + \beta D^{\text{Twin}}_{h,t} + \gamma Z_{h,t} + \varepsilon_{p,h,t}, \tag{R2}$$

for a household h at date $t = \{2002, ..., 2006\}$, where $(\exp^{\text{Educ.}}_{h,p,t}/n_{h,t})$ denotes the education expenditure household h spends on each child (as a share of household income) at date $t = \{2002, ..., 2006\}$.[47]

Results of regression (R2) are shown in columns (2) and (4) of Table 6. 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)).

The quantity–quality trade-off is also visible looking at differences in education attainment. Table 7 displays LOGIT regression results on dummies measuring the level of school enrollment (academic high school, technical high school and higher education). Comparing education attainment of twins versus only children (of age 18–22) over the period 2002–2009 indicates that twins are 40% less likely to pursue

46. If "artificial" twinning was driving our results, then differences between the two types of households would increase over time-"artificial" twinning technologies becoming more accessible. Our investigation does not support this hypothesis. While the saving rate of only child households is higher than twin households, the difference between the two has not risen over time. The average household income is similar between twin and non-twin households (by first child birth) since 1970.

47. Education expenditures are only available for the years 2002–2006 in UHS.

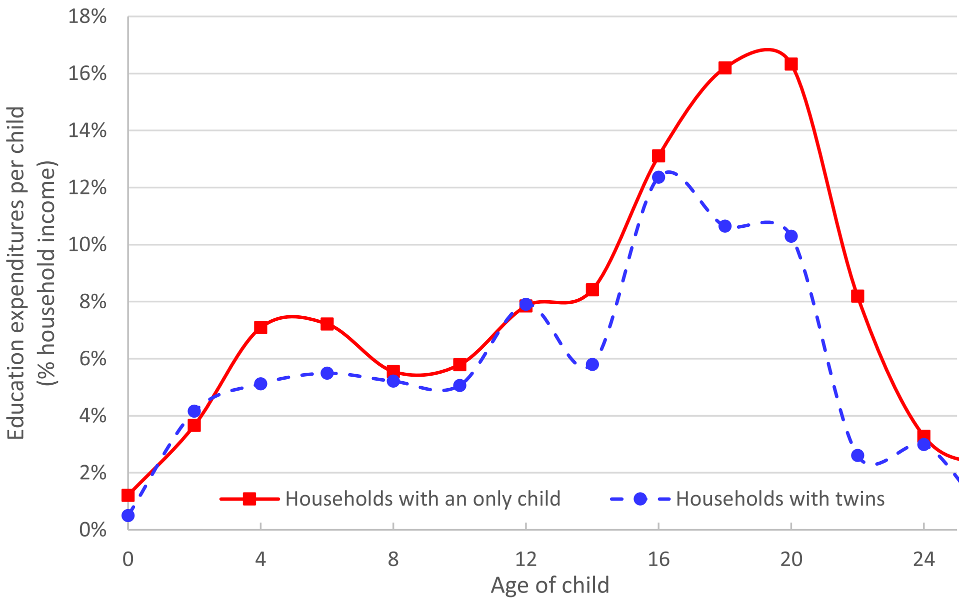

FIGURE 10. Education expenditures per child: Only child versus twins. UHS (2002–2006), restricted sample of nuclear households. This figure displays the average education expenditure per child (as a share of total household income) by age of the child, over the period 2002–2006.

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 preparing to university (Columns (4)) and 30% more likely to attend a technical high school (column (6)).[48]

5.2. *Model versus Data*

Predictions of the "Twin Effect": Model versus 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 empirical estimates. Table 8 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.[49] The predicted saving rates at $\gamma = 9$–10 and $\gamma = 11$–12 are, respectively, 5% (4.9%–5.4% in the data) and 8.0%

48. Twins could be of lower quality compared to singletons—for example, by having lower weights at birth—and parents may in turn invest less in their education. The problem is less serious, however, when households are allowed only one birth as in China. Oliveira (2016) finds no systematic differences between singletons and twins.

49. We estimate the difference across bins of 8 years to preserve a sufficient number of observations for twins.

TABLE 6. Education expenditures per child: twin identification.

Variables (in % of household income)	(1) Education expenditure total	(2) Education expenditure per child	(3) Education expenditure total	(4) Education expenditure per child
Twins	0.0648***	−0.0215***	0.0533***	−0.00917*
	(0.0108)	(0.00539)	(0.0101)	(0.00510)
Twins ≥ 15			0.0277	−0.0248**
			(0.0225)	(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: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 7. Education attainment: twin identification (LOGIT).

	Higher education		Academic high school		Technical high school	
Variable (logistic regression)	(1) Estimate	(2) Odds ratio	(3) Estimate	(4) Odds ratio	(5) Estimate	(6) Odds ratio
Twins	−0.489***	0.613***	−0.455***	0.635***	0.269*	1.308*
	(0.158)	(0.0968)	(0.138)	(0.0875)	(0.157)	(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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

(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)

TABLE 8. Twin experiment: model and data.

	Model			Data[a]
	Only child	Twins	Difference	difference
Saving rate				
$\gamma = 7 - 8$ (25–32)	7.1%	5.8%	1.3%	-1.1 to $2.9\%^{(n.s.)}$
$\gamma = 9 - 10$ (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%	$-0.9\%^{(n.s.)}$
$\gamma = 9 - 10$ (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%	$\left(\frac{H_{\text{Only}} - H_{\text{Twin}}}{H_{\text{Only}}}\right) = 24\%$	

a. Estimates 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)–(4) of Table 4, and the specification for the education expenditures regression is similar to column (1) of Table 6. For the age bracket [49–56], we use the saving estimation based on "movers". $***$ (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).

higher expenditures at $\gamma = 11$–12.[50] 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 is reassuring since the model is not calibrated on twin household variables.

6. Conclusion

We show in this paper that 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

50. In the model, parents of twins thus spend 1.1 percentage points less on education per child (% of wages) at $\gamma = 11$–12.

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 life-cycle model—where human capital investment and intergenerational transfers toward 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.

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 life-cycle 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. 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—giving further credence to the validity of our quantitative model.

This paper demonstrates that shifts in demographics as understood through the lens of a life-cycle 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 becomes older and exerts a greater impact on the economy—both in human capital and in demographic weight. We may therefore expect the effect of the policy on aggregate outcomes to remain in years to come, before the aging of the generation of only child and the progressive relaxation of fertility constraints in China eventually reverse the effects.

References

Allen, Franklin, Qian Jun, and Xian Gu (2015). "China's Financial System: Growth and Risk." Working paper, Imperial College London.

Angrist, Joshua D., Victor Lavy, and Analia Schlosser (2010). "Multiple Experiments for the Causal Link between the Quantity and Quality of Children." *Journal of Labor Economics* 28, 773–824.

Banerjee, Abhijit, Xin Meng, Tommaso Porzio, and Nancy Qian (2014). "Fertility Change and Household Savings in China: Evidence from a General Equilibrium Model and Micro Data." Working paper, Yale University.

Banerjee, Abhijit, Xin Meng, and Nancy Qian (2010). "The Life Cycle Model and Household Savings: Micro Evidence from Urban China."NBER Working paper.

Barro, Robert J. and Gary S. Becker (1989). "Fertility Choice in a Model of Economic Growth." *Econometrica*, 57, 481–501.

Barro, Robert and Jong-Wha Lee (2010). "A New Data Set of Educational Attainment in the World, 1950–2010." *Journal of Development Economics*, 104, 184–198.

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.

Blanchard, Olivier J. and Francesco Giavazzi (2005). “Rebalancing Growth in China: A Three-Handed Approach.” MIT Department of Economics Working Paper No. 05-32.

Boldrin, Michele and Jones Larry (2002). “Mortality, Fertility, and Saving in a Malthusian Economy.” *Review of Economic Dynamics*, 5, 775–814.

Browning, Martin (1992). “Children and Household Economic Behavior.” *Journal of Economic Literature*, 30, 1434–1475.

Bussiere, Matthieu, Yannick Kalantzis, Romain Lafarguette, and Terry Sicular (2013). “Understanding Household Savings in China: The Role of the Housing Market and Borrowing Constraints.” MPRA Paper No. 44611, University Library of Munich.

Cai, Yong (2010). “China’s Below-Replacement Fertility: Government Policy or Socioeconomic Development?.” *Population and Development Review*, 36, 419–440.

Card, David (1999). “The Causal Effect of Education on Earnings.” In *Handbook of Labor Economics*, Vol. 3A, ed. Orley Ashenfelter and David Card, pp. 1801–63. Amsterdam: Elsevier Science.

Chakrabarti, Rupa (1999). “Endogenous Fertility and Growth in a Model with Old Age Support.” *Economic Theory* 13, 393–416.

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

Chesher, Andrew (1998). “Individual Demands from Household Aggregates: Time and Age Variation in the Composition of Diet.” *Journal of Applied Econometrics*, 13, 505–524.

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.

Coeurdacier, Nicolas, Stephane Guibaud, and Keyu Jin (2015). “Credit Constraints and Growth in a Global Economy.” *American Economic Review*, 105(9), 2838–2881.

Coeurdacier, Nicolas, Stephane Guibaud, and Keyu Jin (2014). “Fertility Policies and Social Security Reforms in China.” *IMF Economic Review*, 62, 371–408.

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

De Nardi, Mariacristina, Eric French, and John B. Jones (2010). “Why Do the Elderly Save? The Role of Medical Expenses.” *Journal of Political Economy*, 118, 39–75.

Du, Qingyuan and Shang-Jin Wei (2013). “A Theory of the Competitive Saving Motive.” *Journal of International Economics*, 91, 275–289.

Ge, Suqin, Dennis T. Yang, and Jusen Zhang (2012). “Population Policies, Demographic Structural Changes, and the Chinese Household Saving Puzzle.” IZA Discussion Paper No. 7026.

Gourieroux, Christian S., Alain Monfort, and Alain Trognon (1984). “Pseudo Maximum Like hood Methods: Applications to Poisson Models.” *Econometrica*, 52 , 701–720.

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.

Huang, Wei, Xiaoyan Lei, and Yaohui Zhao (2016). “One-Child Policy and the Rise of Man-Made Twins.” *The Review of Economics and Statistics*, 98, 467–476.

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

Jones, Larry E., Alice Schoonbroodt, and Michele Tertilt (2010). “Fertility Theories: Can they explain the Negative Fertility-Income Relationship?.” *Demography and the Economy*, pp. 43–100, ed. John Shoven, University of Chicago Press.

Jones, Larry E. and Michele Tertilt (2008). “An Economic History of Fertility in the U.S.: 1826–1960.” *Frontiers of Family Economics*, Vol. 1 ed. by Peter Rupert Emerald Press.

Lan, Lan (2019). “Understanding the Saving Behavior of Chinese Households: Intergenerational Transfers and Housing.” Working paper, University of Oslo.

Li, Hongbin, Junsen Zhang, and Yi Zhu (2008). “The Quantity–Quality Trade-Off of Children in a Developing Country: Identification Using Chinese Twins.” *Demography*, 45, 223–243.

Li, Hongbin, Prashant Loyalka, Scott Rozelle, and Binzhen Wu (2017). "Human Capital and China's Future Growth." *The Journal of Economic Perspectives*, 31(1), 25–47.

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.

Manuelli, Rodolfo E. and Ananth Seshadri (2007). "Explaining International Fertility Differences." *Quarterly Journal of Economics*, 124, 771–807.

Modigliani, Franco (1986). "Life-Cycle, Individual Thrift, and the Wealth of Nations." *American Economic Review*, 76(3), 297–313.

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

Newman, John L. (1983). "Economic Analyses of the Spacing of Births." *American Economic Review*, 73(2), 33–37.

Oliveira, Jacqueline (2016). "The Value of Children: Intergenerational Transfers, Fertility, and Human Capital." *Journal of Development Economics*, 120, 1–16.

Psacharopoulos, George and Harry Anthony Patrinos (2018). "Returns to Investment in Education: A Decennial Review of the Global Literature." *Education Economics*, 26, 445–458.

Qian, Nancy (2013). "Quantity–Quality and the One Child Policy: The Only-Child Disadvantage in School Enrollment in Rural China." NBER Working Paper No. 14973.

Raut, Lakshmi K. and T. N. Srinivasan (1994). "Dynamics of endogenous growth." *Economic Theory*, 4, 777–790.

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.

Rosenzweig, Mark R. and Kenneth I. Wolpin (2000). "Natural "Natural Experiments" in Economics." *Journal of Economic Literature*, 38, 827–874.

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.

Rosenzweig, Mark R. and Jusen Zhang (2014). "Co-Residence, Life-Cycle savings and Intergenerational Support in Urban China." NBER Working Paper No. 20057.

Santos Silva, João and Silvana Tenreyro (2006). "The Log of Gravity." *The Review of Economics and Statistics*, 8, 641–658.

Schulhofer-Wohl, Sam (2018). "The Age-Time-Cohort Problem and the Identification of Structural Parameters in Life-Cycle Models." *Quantitative Economics*, 9, 643–658.

Schulhofer-Wohl, Sam and Yang Yang (2006). "APC: Stata Module for Estimating Age-Period-Cohort Effects." Statistical Software Components S456754, Boston College Department of Economics.

Scharping, Thomas (2003). *Birth Control in China 1949–2000: Population Policy and Demographic Development*. Routledge Curzon, London and New York.

Sianesi, Barbara and John van Reenen (2003). "The Returns to Education: Macroeconomics." *Journal of Economic Surveys*, 17, 157–200.

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

Song, Zheng, Storesletten Kjetil, and Fabrizio Zilibotti (2011). "Growing like China." *American Economic Review*, 101(1), 202–241.

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, 1–39.

Storesletten, Kjetil and Fabrizio Zilibotti (2014). "China's Great Convergence and Beyond." *Annual Review of Economics*, 6, 333–362.

Wang, Fei (2012). "Family Planning Policy in China: Measurement and Impact on Fertility." MPRA Paper No. 42226, University Library of Munich.

Wang, Xin and Yi Wen (2012). "Housing Prices and the High Chinese Saving Rate Puzzle." *China Economic Review*, 23, 265–283.

Wan, Junmin (2015). "Household saving and Housing Prices in China." *The World Economy*, 38, 172–192.

Wen, Yi (2011). “Making Sense of China’s Excessive Foreign Reserves.” Federal Reserve Bank of St. Louis Working Paper No. 2011-006A.
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, 511–564.
Yang, Yang, Sam Schulhofer-Wohl, Wenjiang J. Fu, and Kenneth C. Land (2008). “The Intrinsic Estimator for Age-Period-Cohort Analysis: What it is and how to Use it.”1 *American Journal of Sociology*, 113, 1697–1736.
Yang, Dennis T., Junsen Zhang, and Shaojie Zhou (2011). “Why are Saving Rates so High in China?” IZA Discussion Paper No. 5465.

Supplementary Data

Supplementary data are available at *JEEA* online.