

ABSTRACT

Title of Dissertation: ESSAYS ON DEMOGRAPHIC ISSUES IN CHINA
Xue Li, Doctor of Philosophy, 2011

Dissertation Directed By: Professor John Rust
Department of Economics

Chapter 1: Why Is the Sex Ratio Unbalanced in China? The Roles of the One-Child Policy, Underdeveloped Social Insurance, and Parental Expectations

The sex ratio imbalance in China has reached such an alarming level that, by 2020, men of marriageable age are estimated to outnumber women by 24 million. Using a calibrated life-cycle model, this paper examines the rising sex ratio through three linked but different perspectives: one-child policy, social insurance program, and parental expectation. In a dynamic fertility choice framework, a couple's decision on sex selection is motivated by better returns from investing in a son than in a daughter. I also consider the largely overlooked effect of expected sex imbalance on current fertility choices.

The benchmark calibration demonstrates three results. First, moving to a one-and-half-child policy (second allowed if the first is a girl) would dramatically decrease the sex ratio at birth from 125 to 106. Second, if parents are adaptive and take the "cannot-marry" risk into consideration, then the sex ratio under the one-child policy will drop from 125 to 110, while the change in population growth is negligible. Third, when social insurance coverage is universal, the sex ratio only changes by a small amount if parents do not modify their expectation on children's transfer. I also investigate the equilibrium sex ratio when couples are fully rational and forward-looking. If more couples behave in

such a manner, the sex ratio would fall; this suggests that publicity and education could help alleviate the sex imbalance problem in China. In a similar spirit, I consider the issue of endogenizing children's transfer to parents. In an infinite-horizon dynastic model, the equilibrium level of transfer is positively related to the attention parents place on grand-parents' welfare. Finally, I show that if social insurance could change the social attitude on expected child transfer, then it has the potential to significantly reduce the sex ratio.

Chapter 2: Risky Child Investment, Fertility and Social Insurance in China

This paper tries to explain the decline in total fertility rate (TFR) in China by investigating the quantitative effect of social insurance on peoples' fertility choice in an environment where investment in children is risky. The price and income effects of social insurance are heterogeneous depending on peoples' position in the income distribution: low-income people tend to raise more children due to the reinforcing income and price effects, whereas for rich families the income effect dominates the price effect so that their fertility declines in the presence of the social insurance program. Our results based on Chinese economy do not support the hypothesis that increasing social insurance tax rate has a negative impact on fertility rate, as argued in Boldrin, Nardi, and Jones (2005). Through decomposing calibration results under hypothetical policy scenarios and simulating TFRs for various parameter values, we show that liquidity constraints created by a public pension program plays a significant role in *reducing* fertility rate. Factors related to the rate of return on child investment, such as a slowing economic growth, a rise in the cost of childbearing, and potential social attitude changes such as expectations of lower transfers, also contribute to the long-term declining trend in fertility observed in the data.

ESSAYS ON DEMOGRAPHIC ISSUES IN CHINA

by

Xue Li

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2011

Advisory Committee:
Professor John Rust, Chair
Professor Ginger Zhe Jin
Professor Pablo D'Erasmus
Professor Raymond Guiteras
Professor Feinian Chen

© Copyright by
Xue Li
2011

Dedication

To my family.

Acknowledgments

I am deeply indebted to my advisor, John Rust, for his continuous guidance, support and encouragement. His insightful ideas always inspire me to think from various aspects of my topic and his invaluable advice has shaped my work and will continue to influence my future research. I also gratefully acknowledge Ginger Jin for her valuable comments and discussions at each and every stage of the dissertation.

I owe my gratitude to Mico Loretan for his extremely useful comments, and Sandra Ospina for her generous encouragement. Special thanks go to Miranda Moore for generously dedicated her time to reading my draft and giving excellent editorial comments. I also owe my thanks to Pablo D’Erasmus and Raymond Guiteras for their thoughtful questions and suggestions, and Feinian Chen for kindly accepting to serve on my committee. I also benefit from discussions with Daniel Gilligan, Xiaobo Zhang, Macro Cosconati, Jeanne Lafortune, Soohyung Lee, and Yoo-Mi Chin. Finally, I thank participants of the Maryland Applied Microeconomics Seminar, the Eastern Economic Association Annual Meetings (New York), the Midwest Economics Association Annual Meeting (St. Louis), and the North American Summer Meeting of the Econometric Society (St. Louis) for their suggestions and comments.

Table of Contents

List of Tables	vi
List of Figures	vii
1 Why Is the Sex Ratio Unbalanced in China? The Roles of the One-Child Policy, Underdeveloped Social Insurance, and Parental Expectations	1
1.1 Introduction	1
1.2 Background and Significance	6
1.3 The Model	12
1.3.1 Structural Framework under One-Child Policy	14
1.3.2 Structural Framework under Two-Child Policy	22
1.4 Model Results	25
1.4.1 Parameter Choices	25
1.4.2 Benchmark Calibration	27
1.4.3 Heterogeneity with respect to Parents' Income	31
1.4.3.1 The Relationship between Sex Ratio and Parents' Income	32
1.4.3.2 Heterogeneous Impacts of the Three Factors	38
1.4.4 Sensitivity with respect to Sex Selection Price	40
1.4.4.1 Fixed Cost	41
1.4.4.2 Income-Varying Cost	42
1.4.5 Discussion on Model Fitting	43
1.4.5.1 Data	44
1.4.5.2 Contributing Factors in the Cross Section and Over Time	46
1.4.5.3 Model Fitting	48
1.5 Three Extensions	53
1.5.1 Equilibrium Sex Ratio when Parents are Forward-Looking	53
1.5.2 Endogenizing Children's Transfer Rate Distribution	56
1.5.3 A Further Look at the Role of Social Insurance	60
1.6 Conclusion	63
1.7 Appendix	66
1.7.1 Numerical Root-Finding Procedure	66
1.7.2 The Solution for Two-Child Models	69
1.7.3 Benchmark Parameters	71

2	Risky Child Investment, Fertility, and Social Insurance in China	80
2.1	Introduction	80
2.1.1	Motivation	80
2.1.2	International Background	82
2.1.3	Contribution	85
2.2	The Model	86
2.2.1	A Generic Model	87
2.2.2	Does Social Insurance Really Reduce Fertility Rate?	92
2.3	Calibration	95
2.3.1	Parameter Choices	96
2.3.2	Calibration Results	99
2.3.3	Analysis of Calibration Results	100
2.4	Heterogeneity Analysis	104
2.4.1	Total Fertility Rate	105
2.4.2	Average Saving Rate	107
2.4.3	Transfer Rate	108
2.5	Sensitivity Analysis	109
2.5.1	Childbearing Cost	110
2.5.2	Potential Changes in Social Attitudes	112
2.5.3	Social Insurance Tax Rate	113
2.6	Conclusion	117
	Bibliography	131

List of Tables

1.1	Sex Ratios by Ethnic Groups and Socioeconomic Development in 2000 . . .	11
1.2	Calibrated Society-Wide Sex Ratio	29
1.3	Actual versus Simulated Sex Ratio	45
1.4	Actual versus Simulated Sex Outcomes by Sex of Existing Children	51
1.5	Equilibrium Sex Ratio when Parents are Forward-Looking	55
1.6	Benchmark Parameters	72
1.7	Relative Income Distribution for Calibration	73
1.8	Society-Wide Sex Ratio under Different Choices of Fixed Cost c	77
1.9	Society-Wide Sex Ratio under Different Choices of Income-Varying Cost ϕ	78
1.10	Simulated Equilibrium Distribution of Children's Transfer Rate	78
1.11	Sex Ratio Based on Simulated Equilibrium Distribution of Children's Transfer Rate	79
1.12	Calibrated Sex Ratio for Different Channels of Social Insurance	79
2.1	Benchmark Parameters	121
2.2	Relative Income Distribution for Calibration	122
2.3	Calibrated TFR and Private Saving Rate under Hypothetical Environments	122
2.4	Effects of Other Factors on TFR	122

List of Figures

1.1	China Sex Ratio (1976-2008) and Social Insurance Coverage (1993-2008)	9
1.2	Decision Tree under a One-Child Policy	15
1.3	Decision Tree under a One-and-Half-Child Policy	23
1.4	Decision Tree under a Two-Child Policy	23
1.5	Maximal Utility: 1BoySex v.s. 1Girl	33
1.6	Maximal Utility Difference Between 1BoySex and 1Girl Decisions	34
1.7	Maximal Utility: Whether to Have Children	34
1.8	Disproportionate Sex Selection Pattern	35
1.9	Extra Boys from Sex Selection	37
1.10	Sex Ratio by Income	74
1.11	Sex Ratio under Different Fixed Cost c	75
1.12	Sex Ratio under Different Income-Varying Cost ϕ	76
1.13	Numerical Root-Finding: An Example	76
2.1	Optimal Fertility Choice across Parents' Income	123
2.2	Optimal Private Saving Rate across Parents' Income	123
2.3	Optimal Transfer Rate across Parents' Income	124
2.4	Optimal Fertility Choice and Fixed Cost of Childbearing a	124
2.5	Optimal Fertility Choice and Income-varying Cost of Childbearing b	125
2.6	Optimal Fertility Choice and Parents' Altruism Degree towards Grand- parents η	125
2.7	Optimal Transfer/Saving Rate and Parents' Altruism Degree towards Grand- parents η	126
2.8	Optimal Fertility Choice and Expected Child Transfer Rate $\mathbb{E}_t[d_{t+1}]$	126
2.9	Income Effect of an Increase in α_t	127
2.10	Optimal Fertility Choice and Social Insurance Tax Rate α_t	128
2.11	Optimal Fertility Choice and Expected Social Insurance Tax Rate α_{t+1}	128

Chapter 1

Why Is the Sex Ratio Unbalanced in China? The Roles of the One-Child Policy, Underdeveloped Social Insurance, and Parental Expectations

1.1 Introduction

*“When a son is born,
Let him sleep on the bed,
Clothe him with fine clothes,
And give him jade to play...
When a daughter is born,
Let her sleep on the ground,
Wrap her in common wrappings,
And give broken tiles to play...”*

— Taken from China’s *Book of Songs*¹ (1100-600 B.C.)

China, with a traditional preference for boys, faces growing gender imbalance among newborns since the 1980s: the national average was 119 in 2005, far exceeding the United Nations’ recommendations (no more than 107); and a significant number of cities had sex ratios higher than 125. An official estimate from China Academy of Social

¹The *Book of Songs* is the earliest existing collection of Chinese poems and songs. It is regarded as a revered Confucian classic, and has been studied and memorized by centuries of scholars in China. The above excerpts are from the 189th poem titled “Si Gan”, which is recorded on Chapter 4 (Decade of Qi Fu), Section II (XiaoYa, or Minor Odes of the Kingdom) of the *Book of Songs*.

Sciences suggests that, by 2020, there may be 24 million men of marriageable age who will not be able to find a wife. This vast army of surplus males could lead to social instability; for example, human trafficking and forced prostitution have become “rampant” in some parts of the country. With the ticking time bomb of the sex ratio imbalance, one would ask how did this come about, and how can China address this problem?

Son preference, as vividly shown in the ancient poem above, is always considered as the root cause for gender imbalance. It comes from two sources: one is the cultural aspect — males could carry family names and inherit family properties; the other is the economic aspect — sons could provide more old-age support than daughters. Despite this well known preference for sons that has existed in China for thousands of years, the serious gender imbalance is only a relatively recent phenomenon. Indeed, even as recently as the three decade interval from 1950 to 1980, sex ratios were only slightly higher than the ratio of 105 male births per 100 female births that is considered to be a normal sex ratio due to a variety of biological factors², and is observed in many other countries that do not have the strong cultural bias in favor of sons that China and several other Asian countries have³. However, the sex ratio started to soar in the 1980s. The most intuitive explanation is the enforcement of the one-child policy⁴, which induces widespread abortions on female fetuses. Then our first question is: how much does this

²In reality, the natural sex ratio at birth is between 103 and 107 (on average 105), indicating a slightly higher probability of having a son, where the higher probability of having a baby boy is used to compensate the higher infant mortality risk for males so that the sex ratio evens out in adult population.

³In the 1953 China population census, sex ratio at age 0 was 104.9 boys per 100 girls. According to Ma et al. (1998), average sex ratio at age 0 was 106.9 in the 1950s, 107.5 between 1960 and 1969, and 106.0 in the 1970s.

⁴China’s one-child policy was established in 1978 and the enforcement remains strong as of 2008. However, this birth control policy is a diversified program in that although urban residents can only have one child; many rural couples are allowed a second kid if the first is a girl; ethnic minority couples are allowed to have two or more children; and no restrictions in Tibet. The latest revision is that couples in which both partners are single children may be allowed to have two.

state-mandated family control policy contribute to the sex imbalance?

With respect to motives for childbearing, old-age support is often mentioned, and social insurance program is cited to help reduce both population growth and sex imbalance. In particular, China has made some progress in the development of its social insurance system since 1993; however, the increasing trend in the sex ratio did not stop or slow down after China officially launched this program. So, our second question is: what is the role of social insurance in a couple's fertility decision and how does it affect the society-wide sex ratio?

Last but not least, the recent trends in the sex ratio since 2000 seem to indicate a new pattern, i.e. the sex ratio has remained at alarmingly high level, between 120 and 125. Numerous news reports⁵, both in China and abroad, have expressed concerns that with such a high imbalance, marriage markets in 20 years will be extremely unfavorable to boys. Then this brings about our third question: is parental expectation in terms of concerns on the marriage prospect for sons able to reduce the sex ratio? This is a challenging issue since it involves policy debates on whether China should reform its family planning policy, and social-economic concerns on whether such a high sex ratio is permanent or transitory. Here, we consider the possibility that people recognize the "can-not-marry" risk for sons and expect lower support from unmarried sons, and we use our model to study how this parental expectation change would affect fertility choices and the sex ratio.

To answer all these questions, one needs to understand how much of the current sex

⁵For example, in July 2010, People's Daily (the official newspaper of China's central government) had an article titled "Brides for Sales: Sex Ratio Imbalance Troubles China" discussing a series of problems related to the severe gender imbalance among young Chinese. BBC News had a similar article featuring "Wifeless Future for China's Men" as early as in February 2007.

imbalance is due to the one-child policy; if relaxation could help alleviate the problem, to what extent the policy should be relaxed; and the direction and magnitude of social insurance's impact. What's more, given the pros and cons of the above two, whether changes in parental expectation could be helpful in easing the unbalanced gender structure. All these involve a great deal of variations in the policy environments couples face when making fertility choices. Since it is hard to impose nationwide experiments to determine whether a reform is effective or not, we calibrate an individual decision making model to address this issue from various aspects and to shed some light on these intriguing questions.

We develop a tractable life-cycle model that captures couples' decisions regarding (1) whether or not to have children; (2) if the sex of a fetus is a girl, whether to abort in order to try again for a boy; (3) depending on the birth quota, whether to have a second child; and (4) if a second kid is allowed, whether to terminate a girl pregnancy on the second child as well. Along with these fertility choices, couples also make optimal decisions on consumption, transfer to their elderly parents, and personal savings. One major feature of our model is that investment in children is costly (measured by money and time), while the return on this investment is subject to several risks (child mortality risk, adult "can-not-marry" risk, and adult transfer uncertainty). These decisions are modeled encompassing three different dimensions: (1) whether the one-child policy is enforced, (2) whether social insurance coverage is available, and (3) whether parents are forward-looking with respect to their sons' marriage prospects. The third dimension in our model is the key ingredient of the reverse effect of the sex ratio on fertility choices: if couples rationally expect that a high sex imbalance will dash their sons' chances of marriage in the future and reduce their expected transfers from unmarried sons, then their preferences

and choices over sons and daughters may change.

Our integrated model allows us to investigate the potential causes of the gender imbalance problem and to understand possible outcomes from different policy experiments. As predicted, moving from a stringent one-child policy to a one-and-half-child policy (second allowed if the first is a girl) would dramatically decrease the sex ratio at birth. The impact of social insurance is more complicated in that it involves four different channels (income effect, price effect, liquidity constraints, and social attitude changes). At this moment without the change in social attitudes, its overall magnitude is limited as compared to that of the family control policy. However, when parents are forward-looking and take into account the “can-not-marry” risk for sons, the sex ratio declines significantly without a noticeable increase in the total fertility rate. This suggests that changes in parental expectation may alleviate the sex imbalance problem and simultaneously avoid a higher population growth, concerns over which are precisely why the Chinese government are resistant to reforming the controversial family planning policy.

The rest of the paper is organized as follows. Section 2 provides background information on the one-child policy and the social insurance program, and discuss the significance of this sex ratio imbalance issue for China. Section 3 describes the dynamic fertility choice framework with utility maximization for couples. Section 4 presents the main model results and evaluates the model’s goodness of fit by comparing actual versus simulated sex ratios. We illustrate how the one-child policy, the social insurance program, and parental expectation affect agents heterogeneously. And we examine how our model can be used to evaluate counterfactual policy experiments. For example, we show that an increase in the sex selection cost such as strengthening the supervision on non-medical

abortions would result in a significantly lower sex ratio. Section 5 provides three extensions to the benchmark framework. First, we consider the scenario when parents are fully rational and forward-looking, and compare the steady state to the current sex ratio. Second, we endogenize children's transfer behavior by deriving an equilibrium transfer distribution in an infinite-horizon framework. Finally, we discuss additional channels through which social insurance can affect fertility decisions and show that if the introduction of social insurance program changes the social attitudes on child transfer, then it could significantly affect the sex ratio. We offer concluding remarks in Section 6.

1.2 Background and Significance

This study is inspired by Sen (1990), who draws attention to an important fact of life in East and South Asia: a biased sex ratio at birth and males outnumbering females. In China, India and South Korea, gender imbalance has become a longstanding problem due to various human interventions, from sex selective abortions to neglect or even infanticide as seen in the substantial female child mortality as discussed, for example, in the March 6, 2010 issue of the *Economist* magazine. The fundamental reason for this sex ratio imbalance is the persistent son preference⁶ in these countries. There are two separate, though not independent, causes for this preference. First, these countries share strong similarities in their rigidly patrilineal kinship system, which lies at the root of discrimination against daughters (see Das Gupta et al. (2003)). Second, economic factors including old-age support, dowries, labor force participation, etc., may account, to various degrees, for the son

⁶Actually, Williamson (1976) argued that most societies show some degree of preference for sons, though most are so mild as to be virtually undetectable.

preference⁷.

Among all these countries, China deserves a special attention: it is the country with the biggest population, and it faces the most severe gender imbalance. More importantly, the imbalance of sex structure recorded in the past three decades is not only a demographic problem, but also an issue affecting every aspect of the society such as population size, aging, a wifeless future and social stability. Under the state-mandated one-child policy, China's total fertility rate remains low resulting in an increasing proportion of elderly people in the society. Meanwhile, the abnormally high sex ratio will lead to a "marriage squeeze" for young adult males, with predictions of as many as 24 million men of marriageable age not able to find a wife by 2020. As argued in Wei and Zhang (2011), in order to increase the attractiveness of sons in the marriage market, Chinese parents have strong biological motivation to save, which may contribute to the high saving rate and the soaring housing prices in China. Moreover, these surplus males often play a crucial role in making violence prevalent within the society and thus harm social stability. Edlund, et al. (2007) document the relationship between sex imbalance and the increase in the crime rate in China.

Previous literature focuses on two major aspects in explaining China's gender imbalance: family planning regulation associated with sex-selective abortions, and underdeveloped social insurance. First, the one-child policy narrows peoples' fertility choice set and stimulates couples to find ways to satisfy their son preference. Ultrasound technology provides a means to do sex-selective abortions at a reasonable cost. Using a difference-

⁷For example, Qian (2008) claims that increasing female income, holding male income constant, improves survival rates for girls and increases educational attainment of all children. Rao (1993) and Anderson (2003; 2007) discuss the inflation of dowry payments, brideprice, and female power in India.

in-differences method, Li et al. (2010) conclude that the one-child policy has resulted in around 7.0 extra boys per 100 girls for the 1991-2005 birth cohort and accounts for between 54% and 57% of the total increase in sex ratio for the 1990s and the 2001-2005 birth cohorts. On a related subject, Li and Zheng (2009) try to directly measure the causal effect of sex selective abortions on the sex ratio at birth by exploiting the exogenous county-level variation in the availability of ultrasound machines. They find that such availability increases the sex ratio at birth by 0.025 in rural and 0.117 in urban areas. Second, social insurance could arguably ameliorate the differing old-age support from sons and daughters: a generous pension benefit could substitute part of a son's role⁸. For example, Ebenstein and Leung (2010) show that people who have sons are less likely to enroll in voluntary social insurance program, and the sex ratio is mitigated in counties with old-age pension programs. By the same token, Bhattacharjya, et al. (2008) argue that policies involving economic benefit (such as pension plans for families with no sons) could decrease the difference between the perceived present value of sons and daughters, and thereby reduce the sex ratio.

Few of these empirical studies have taken an integrated structural approach to consider different factors simultaneously. In addition, the difference-in-differences and "treatment effects" methods used in any reduced form studies cannot accurately reflect the complexity and uncertainty facing heterogeneous individual decision makers; nor do they capture the critical dynamic elements of fertility choices. Moreover, it is very hard to predict how new, hypothetical policy changes might affect outcomes in the future us-

⁸The sex ratio in South Korea reached its peak value of 117 in the 1990s. As of 2008, it had dropped to a close-to-normal level of 107 and anecdotal evidence indicates that a series of reforms to social security was partially responsible for this drop in sex ratios.

ing a reduced form methodology. In providing guidance for policy makers, it is critical to be able to predict the consequences of hypothetical counterfactual policy experiments. Therefore, in this paper, we will apply a structural framework to analyze individual optimal choices and forecast their responses to a wide range of policy changes, such as relaxing the one-child policy, strengthening regulations on sex-selective abortions, promoting social insurance to rural areas, and educating the general public that girls are equally good as boys.

Figure 1.1: China Sex Ratio (1976-2008) and Social Insurance Coverage (1993-2008)




Figure 1.1 describes China’s sex ratio history since 1976. Clearly, sex ratio began to increase after the enforcement of the family control policy in the late 1970s; this increasing tendency seems to have halted recently, but the sex ratio for the age 0-4 group remains around 120 with some fluctuations; and the future trend appears unclear at this moment⁹. We also display the social insurance coverage rate from 1993 onwards, which

⁹Das Gupta et al. (2009) argues that in China, the provinces which had the highest sex ratios (and have two-thirds of China’s population) have seen a deceleration in their ratios since 2000, and provinces with

is calculated as the number of people having old-age pension coverage¹⁰ divided by the total population aged 15 and above. Although social insurance is still underdeveloped in China, people covered under this system almost doubled from 1993 to 2008. However, the increases in the sex ratio and in the social insurance coverage rate seem to be parallel to each other, which could suggest that social insurance may not have a significant impact on the sex imbalance.

We also look at the sex ratio by ethnic groups and socioeconomic development for year 2000 in Table 1.1. A comparison between Han Chinese and other ethnic groups gives us a rough idea of the effect of the one-child policy. Han Chinese, accounting for over 90% population, had higher sex ratios (119 nationwide) than the minority groups (112 nationwide), who are exempt from the one-child regulation. Another clear observation is that sex ratio was lower in cities than in towns and villages. But the reasons behind this phenomenon are not that apparent. One possibility is that cities have relatively better developed social insurance programs. However, if this were true, we should observe (in Figure 1.1) a slowing in the increase in sex ratio since 1993 when the program was introduced; but we do not. This suggests that lack of social insurance coverage may not be the main force pushing up sex ratios in rural areas. In addition to social insurance, rural and urban areas are different in several other aspects as well, which could potentially contribute to the observed differences in sex ratios. Abortion is cheaper and regulation of illegal sex-selective abortions is weaker in rural areas than in urban areas. Young adult

a quarter of the population have seen their ratios fall. This, at the very least, seems to be an incipient turnaround of the “missing girls” phenomenon.

¹⁰China’s social insurance system has four parts: the old-age pension program, health insurance, unemployment insurance, and maternity insurance. The latter three are less developed than the first one and cover much fewer people. Therefore, our calculation should give an upper bound for social insurance coverage rate.

males may be more valuable for farm work than females and they have a higher potential income (for instance, they can migrate to a city), enlarging the difference in rewards between sons and daughters for rural families. Housing prices in cities are less affordable and it is a social custom for a bride's parents to buy a house for the marriage, decreasing the parents' motivation for sex selection, etc.

Table 1.1: Sex Ratios by Ethnic Groups and Socioeconomic Development in 2000

	Nationwide	City	Town	Village
Han Chinese (91.9% of total population)	119	113	118	120
Other 55 Ethnic Groups (8.1% of total population)	112	108	112	113

Note: Sex ratio (boys/girls) at age 0 is calculated using the summary statistics of population by age, sex, and nationality of the 2000 population census.

Source: 2003 China Population Statistics Yearbook.

Inspired by the potential effect of housing prices on fertility choices and numerous news reports on the millions of surplus males facing a wifeless future, we investigate a largely overlooked area: the (reverse) effect of the sex imbalance on couples' fertility choices. One can imagine a marriage market in 20 years, in which some males can not find wives because of the sex imbalance. Their transfers to their parents may be lower than if they can marry, which suggests that daughters should be at a premium and bring more rewards to parents. If Chinese parents are aware of the environment of excess boys and treat it in a serious manner, they should rationally react to the current sex ratio so that the high imbalance would only be a short-term phenomenon. However, this has not happened in reality and a partial justification could be that Chinese parents haven't realized the ensuing tight marriage market for sons, or some couples may have biased interpretations (for example, they may be over-confident of their son's chance of marriage). To

consider this channel, we incorporate explicitly how parental expectation on future transfers are formed via an individual optimization model. We will consider three types of expectations: myopic, adaptive, and completely forward-looking, and see how parental expectations move the sex ratio.

1.3 The Model

The key feature of our theoretical framework is a macro aggregation based on micro optimization. Specifically, we consider agents with heterogeneous budget constraints, solve their individual optimization in a partial equilibrium framework, then aggregate individual fertility choices to obtain the society-wide sex ratio. This is different from a typical representative agent problem in a general equilibrium framework such as Boldrin and Jones (2002). In such a general equilibrium framework, the rule for optimal behavior is the same across periods and there is no uncertainty in child transfer¹¹ from the parents' viewpoint. In addition, since the representative agent is assumed to represent the whole society, his decisions could affect the society outcome, with wage and interest rate endogenously determined. However, our framework is micro-founded in that agents in this framework are heterogeneous, since we do not assume that the aggregate behavior of

¹¹There is an abundant literature studying the direction of intergenerational family transfers, the underlying motives, and the supporting institutional and cultural arrangements, represented by Caldwell (1978; 1982), Willis (1982), etc. Consistent with Caldwell's "old-age security" hypothesis, Boldrin and Jones (2002) model the fertility choices that children are investment goods to parents and the desired number of children depends on the amount the child transfers to elderly parents in relation to the cost of rearing their child to adulthood. This is contrary to the work of Barro and Becker (1988; 1989) in which the utility of children enters directly into the utility function of the parents, indicating the reason for childbearing is that children are viewed as life continuity for parents. Although we focus on the choice between sons and daughters instead of the optimal number of children, we adopt a framework similar to Boldrin and Jones (2002) to reflect the son preference as better returns from investing in a son than in a daughter in terms of the expected future transfers.

millions of heterogeneous families can be well approximated by the behavior of a single ‘representative consumer’. The population outcomes are from a direct aggregation of individual choices, where we assume that wage rate and interest rate are exogenously given. More importantly, in this partial equilibrium framework, parents are uncertain about the old-age transfers they can expect to receive from their children, so they need to form expectations on how they will be distributed.

Instead of focusing on endogenous economic growth and fertility transitions, I calibrate the model based on optimal decisions solved in a partial equilibrium framework to emphasize the heterogeneous feature of peoples’ reactions to homogenous policy environments (either social insurance or the family planning policy). Under this framework, an individual’s decision could not have impacts on the aggregate outcomes, but the aggregation of individual decisions can affect individual choices via individual expectations. Thus, this is the advantage of introducing heterogeneous agents in the micro-aggregation, and it may produce more realistic results.

In detail, we present a dynamic fertility choice model in a dynamic programming framework. The main fertility decisions are: (1) whether to have child(ren) or not, (2) if have children and if the screening shows it is a girl, whether to abort or not, (3) depending on the birth quota defined by the family control policy, whether to have a second kid, and so on. Payoff at each terminal decision node is determined by an individual utility maximization on consumption, transfer and saving, with the fertility choice at that node taken as given.

Although Ebenstein (2011) presents a similar fertility choice framework¹², there

¹²Ebenstein (2011) assumes parents have access to a priced sex selection technology, but they have to

are some notable differences: (1) we endogenize the payoff of having children by solving a three-period life-cycle model, covering individual's investment in children, transfer to elderly parents, consumption and saving; (2) we introduce social insurance in a way that alters peoples' expectation on the return from child investment, as well as affecting their intertemporal budget constraints; (3) we consider three versions of the family planning policy: one-child, one-and-half-child (second allowed if the first is a girl), and full-fledged two-child policy; (4) aside from the normally considered mortality risk in child investment, we incorporate a "can-not-marry" risk for sons reflecting the reverse effect of the sex imbalance on fertility choices; (5) the heterogeneity in fertility choices with respect to income is assessed, so the distortion of the sex ratio is different depending on peoples' income position. We also allow the "double-income-no-kid" (DINK) phenomenon, which may be optimal for certain type of individuals given their preference and budget constraints.

Overall, our structural model spans three dimensions: variations in the family control policy; the presence of social insurance; and incorporating a son's "can-not-marry" risk into the parents' child investment consideration.

1.3.1 Structural Framework under One-Child Policy

The dynamic decision making process is presented as a decision tree in Figure 1.2.

Using backward induction, we solve the maximization problem for the following scenarios: (1) one-boy with sex selection; (2) one-girl; (3) one-boy, (4) no-children. By

pay a fine in order to have a second birth. He also considers a three-child decision problem to explore the possible outcomes by relaxing the one-child policy.

comparing the expected life-time utility under different scenarios, parents make optimal fertility choices as well as optimal consumption, transfer and saving choices at each decision node. We solve this individual decision making problem for everyone in the society, and then aggregate individual fertility choices into a society-wide sex ratio.

To simplify our calculation, we assume parents have equal probability (i.e. 50%) of having a boy or a girl so that if nobody choose to do sex selection, the sex ratio will be 100. However, given that the natural sex ratio at birth is between 103 and 107 (on average 105), we must be careful in explaining our simulated sex ratio and realize that we need to adjust our results upwards¹³ in order to compare with actual data.

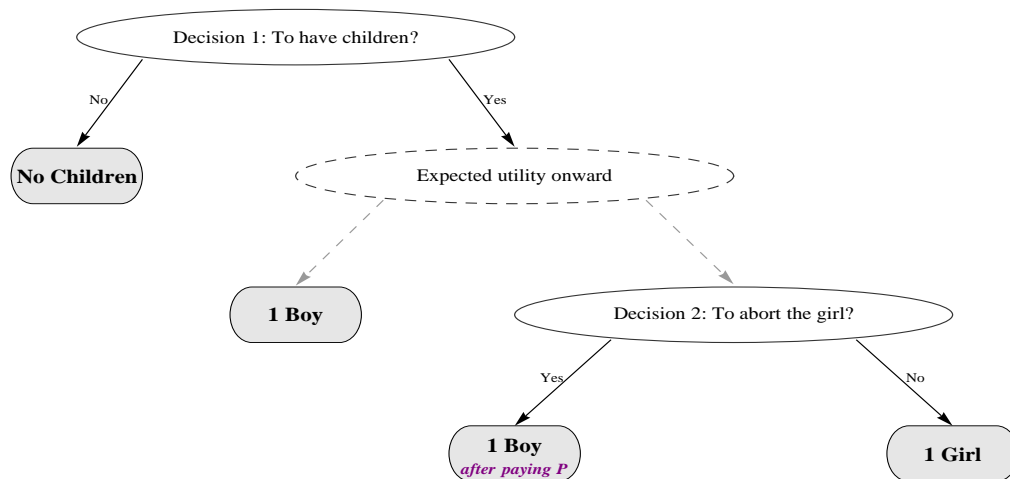


Figure 1.2: Decision Tree under a One-Child Policy

According to Figure 1.2, we implicitly assume that the enforcement of the one-child policy is perfect so that if the incoming baby is a girl, parents can only choose to

¹³There are two ways to adjust. Since the natural sex ratio at birth in our framework is 100, while it is 105 (on average) in reality, the absolute difference is 5 and the relative difference is 5%. If our simulated sex ratio equals 120, one way is to add 5 directly which arrives at 125; the other way is to increase our results by 5% which is 126(=120 × 1.05). We adopt the second way when measuring the model's goodness of fit in section 1.4.5.

do sex selection to try to have a boy, but can not or are not allowed to pay fines to have a second (or even a third) child. We admit that this assumption seems to be away from reality and we do observe families pay fines for violating the policy in order to have more than one kid. However, considering imperfect enforcement of the family control policy will complicate our dynamic fertility choice process and make it hard to have a clear understanding of the effect of the family control policy on sex ratio. Suppose the one-child policy is not perfectly enforced so that parents can pay fines to have a second kid, then the difference in sex ratio between one-child and one-and-half-child scenarios does not indicate the exact impact of relaxing the one-child policy; instead it reflects a combined effect of imposing fines and relaxing the policy. In reality, if the local authority has weak enforcement penalty for violating the policy, we could imagine that the magnitude of relaxing the one-child policy will be significantly under-estimated by comparing sex ratios across different areas. Thus, as a simplification, we assume that the family control policy is perfectly enforced.

On life-cycle dynamics, we assume that individuals can live three periods: young, middle-age, and old. Young individuals simply consume parents' resources to grow up. When becoming middle-aged, they supply one unit of labor, obtain income (W_t) ¹⁴ and make optimal decisions on fertility, consumption (C_t^m) , transfer to their elderly parents (d_t) and saving (s_t) .

The cost of rearing children consists of two parts: a fixed cost (a) and an income-

¹⁴Here is the notation rule: lower-case letters usually represent the percentage or ratio, while upper-case represent the absolute level. For example, \overline{W}_t is the society-wide average income level, $W_{t,i}$ is the income for individual i , and the corresponding relative income is $w_{t,i} = W_{t,i}/\overline{W}_t$. Similarly, $s_{t,i}$ represents the private saving rate for individual i , while $S_{t,i} = s_{t,i}W_{t,i}$ is his/her private saving amount. For ease of notation, we omit the subscript i .

varying cost. When there is no social insurance coverage, the time cost of rearing one child is simply bW_t ; when social insurance is present, it will be discounted by tax rate (α_t), which becomes $(1 - \alpha_t)bW_t$. That is, the opportunity cost of raising a child instead of working decreases because people need to pay social insurance taxes. Correspondingly, the income by providing one unit of labor drops from W_t to $(1 - \alpha_t)W_t$. In a similar spirit, the cost of sex selection also contains two parts: a fixed cost (c) and an income-varying cost (ϕW_t). Here, we assume that the child-care expenses are tax-deductible, but the cost of sex selection is not. On expected transfers next period (D_{t+1}), we decompose it as: $d_{t+1}w_{t+1}\bar{W}_{t+1}$, where we assume the transfer rate d_{t+1} follows a Beta distribution¹⁵ and children's relative income w_{t+1} is log-normally distributed¹⁶. When individuals become old, they will retire and finance their consumption by the transfer from children, private savings from previous period¹⁷ and social insurance benefits if they are covered by the program.

We first present a representative scenario, and then show that other scenarios can be accommodated as special cases. The typical model of “One-boy with sex selection, social insurance and adaptive parents” scenario is:

¹⁵Beta distribution is a family of continuous probability distributions defined on the interval $(0, 1)$ and parameterized by two positive shape parameters, typically denoted by α and β . This Beta distribution assumption guarantees that the transfer rate is always within the $(0, 1)$ region.

¹⁶This assumption ensures that the relative income is positive, and the median income is smaller than the mean, as consistent with empirical income distributions.

¹⁷ f_{t-1} and s_{t-1} are two state variables, indicating the number of children and private saving rate of the middle-aged in the previous period.

$$\text{Max}_{\{s_t, d_t\}} \mathbb{U}(C_t^m, C_{t+1}^o, C_t^o) = \{u(C_t^m) + \delta \mathbb{E}_t u(C_{t+1}^o) + \eta u(C_t^o)\} \quad (1.1)$$

s.t.

$$C_t^m = (1 - s_t - d_t)(1 - \alpha_t)W_t - (a + b(1 - \alpha_t)W_t) - (c + \phi W_t) \quad (1.2)$$

$$C_{t+1}^o = d_{t+1}w_{t+1}(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + \text{SI}_t \quad (1.3)$$

$$C_t^o = f_{t-1}d_t(1 - \alpha_t)W_t + R_t s_{t-1}(1 - \alpha_{t-1})W_{t-1} + \text{SI}_{t-1} \quad (1.4)$$

$$\text{SI}_t = (\gamma_t \bar{W}_t + \beta_t \alpha_t (1 - b)W_t) R_{t+1} \quad (1.5)$$

$$\text{SI}_{t-1} = (\gamma_{t-1} \bar{W}_{t-1} + \beta_{t-1} \alpha_{t-1} (1 - f_{t-1}b)W_{t-1}) R_t \quad (1.6)$$

with F.O.C.

$$u'(C_t^m) = \delta R_{t+1} \mathbb{E}_t (u'(C_{t+1}^o)) \quad (1.7)$$

$$u'(C_t^m) = \eta f_{t-1} u'(C_t^o) \quad (1.8)$$

Assuming a logarithmic utility function and after algebraic manipulations of eq. (1.7) and (1.8), we can write d_t as a function of s_t ,

$$d_t = \frac{\eta}{1 + \eta} \left((1 - s_t - b) - \frac{\phi}{1 - \alpha_t} - \frac{a + c}{(1 - \alpha_t)W_t} \right) - \frac{R_t s_{t-1} (1 - \alpha_{t-1})W_{t-1} + \text{SI}_{t-1}}{(1 + \eta) f_{t-1} (1 - \alpha_t)W_t} \quad (1.9)$$

Thus, the solution is characterized by an equation on s_t , $\text{LHS}(s_t) = \text{RHS}(s_t)$, and we use Newton's method to find its root.

$$\text{LHS}(s_t) = \frac{1 + \eta}{(1 - \alpha_t)(1 - s_t - b)W_t - \phi W_t - (a + c) + (R_t s_{t-1}(1 - \alpha_{t-1})W_{t-1} + \mathbf{SI}_{t-1}) / f_{t-1}} \quad (1.10)$$

$$\begin{aligned} \text{RHS}(s_t) &= \delta R_{t+1} \mathbb{E}_t \left(\frac{1}{d_{t+1} w_{t+1} (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \mathbf{SI}_t} \right) \\ &= \frac{\delta p R_{t+1}}{s_t (1 - \alpha_t) W_t + \mathbf{SI}_t} \\ &\quad + \int_0^\infty \int_0^1 \frac{\delta (1 - p) \mathfrak{m} R_{t+1}}{d_{t+1}^y w_{t+1} (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \mathbf{SI}_t} h_D(d^y) g_W(w) \, ddw \\ &\quad + \int_0^\infty \int_0^1 \frac{\delta (1 - p) (1 - \mathfrak{m}) R_{t+1}}{d_{t+1}^x w_{t+1} (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \mathbf{SI}_t} h_D(d^x) g_W(w) \, ddw \end{aligned} \quad (1.11)$$

where

$$d_{t+1}^x \sim \text{Beta}(\alpha^x, \beta^x) \quad , \quad d_{t+1}^y \sim \text{Beta}(\alpha^y, \beta^y)$$

$$\mathbb{E}_t[d_{t+1}^y] = \lambda \mathbb{E}_t[d_{t+1}^x] \quad , \quad \text{Var}(d_{t+1}^y) = \lambda^2 \text{Var}(d_{t+1}^x) \quad , \quad 0 < \lambda < 1$$

$$\log(w_{t+1}) \sim \mathcal{N}(\mu, \sigma^2)$$

The RHS could be simplified as follows if we construct a discrete approximation to the product of Beta and lognormal distributions¹⁸.

¹⁸Given two independent random variables X and Y , the distribution of $Z = XY$ is a product distribution, whose density can be derived as follows:

$$p_Z(z) = \int_{-\infty}^{+\infty} \frac{1}{|x|} p_{X,Y}\left(x, \frac{z}{x}\right) dx$$

where $p_{X,Y}(x, y)$ is the joint probability density function. Here, d_{t+1} and w_{t+1} are assumed to be statistically independent so $p_{X,Y}(x, y) = p_X(x)p_Y(y)$. Hence, we can derive the distribution of $d_{t+1}w_{t+1}$ and then construct a discrete approximation to this product distribution. Details are provided in Appendix 1.7.1.

$$\begin{aligned}
\text{RHS}(s_t) &= \delta R_{t+1} \mathbb{E}_t \left(\frac{1}{d_{t+1} w_{t+1} (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \mathbf{SI}_t} \right) \\
&= \frac{\delta p R_{t+1}}{s_t (1 - \alpha_t) W_t + \mathbf{SI}_t} + \int_{\mathbb{R}} \frac{\delta (1 - p) \mathfrak{K} R_{t+1}}{e^y (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \mathbf{SI}_t} f_Y(y) dy \\
&\quad + \int_{\mathbb{R}} \frac{\delta (1 - p) (1 - \mathfrak{K}) R_{t+1}}{e^x (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \mathbf{SI}_t} f_X(x) dx \tag{1.12}
\end{aligned}$$

where

$$x \equiv \text{Log}(d_{t+1}^x w_{t+1}) \quad , \quad y \equiv \text{Log}(d_{t+1}^y w_{t+1})$$

$$\mathbb{R} \equiv (-\infty, +\infty)$$

The key source of risk in this model is the intergenerational transfer amount, which is decomposed into two factors: the transfer rate (which is affected by the social attitudes, etc, but will also be some kind of optimal choice when the kids grow up) and the relative income (which is related to parents' human capital investment). We assume that boys and girls differ in the distribution of their transfer rate, but not in that of their relative incomes¹⁹ (i.e. $d_{t+1,\text{boy}} \sim \text{Beta}(\alpha_b, \beta_b)$ and $d_{t+1,\text{girl}} \sim \text{Beta}(\alpha_g, \beta_g)$). In addition, two other risks are also involved: (1) mortality risk denoted as p , which is introduced to ensure that parents driven by precautionary saving motives, always have a positive net asset; (2) “can-not-marry” risk, which is to model different types of parental expectations, such as myopic, adaptive, and completely forward-looking. Here we denote the probability of “can-not-marry” as \mathfrak{K} , which is positively related to the cohort-wide sex ratio; if a son cannot get married, we assume that his income will not be affected by his marriage status,

¹⁹To avoid further complication, we omit the gender discrimination in the labor market.

but he will transfer, on average, a smaller percentage of his income to his elderly parents. That is, unmarried sons' transfer rate still follows a Beta distribution, with the mean being $\lambda \mathbb{E}_t[d_{t+1, \text{boy}}]$ and variance being $\lambda^2 \text{Var}(d_{t+1, \text{boy}})$, where $0 < \lambda < 1$; $\mathbb{E}_t[d_{t+1, \text{boy}}]$ and $\text{Var}(d_{t+1, \text{boy}})$ are the corresponding mean and variance of the transfer rate distribution for those married sons.

Let's look at the parental expectation scenarios further: the completely myopic (or over-confident) scenario implies that $\varkappa = 0$; the adaptive scenario implies that \varkappa is derived from the current sex ratio as $\varkappa = \rho(1 - \frac{100}{\kappa})$, where sex ratio κ is denoted as number of boys per 100 girls and ρ represents some adjustment accounting for cross-cohort marriage, immigration, emigration, etc; and completely forward-looking means there should be an equilibrium level κ^* such that the sex ratio from parents optimally choosing fertility matches with the ex ante expected sex ratio in the mind of these parents. Details on deriving the equilibrium sex ratio are provided in section 1.5.1.

Introducing social insurance does not change much of the model. Social insurance serves as a mandatory saving mechanism, with the rate of return depending on individual's income position; given that the Chinese income distribution is highly skewed to the right (median smaller than mean), it is expected that most people may benefit from this program. Moreover since part of the childcare-related expenses is tax-exempt, everyone should see their childbearing cost (i.e. the income-varying part) lower than that without social insurance. So the income and price effects mean that social insurance should induce most people (except for those with very high income) to have more kids, but its effect on sex ratio is uncertain.

Finally, let's look at some sub-models: (1) "One-girl": we need to set $(c, \phi) =$

$(0, 0)$, and pick up the corresponding (α_g, β_g) ; (2) “One-boy”: we need to set $(c, \phi) = (0, 0)$; (3) “No-kid”: we need to set $(a, b, p) = (0, 0, 1)$ and $(c, \phi) = (0, 0)$, and an analytical solution happens to exist.

1.3.2 Structural Framework under Two-Child Policy

In some areas of China parents can have a second kid only if the first one is a girl, or if they face an economic hardship (which is termed as one-and-half-child policy, as shown in Figure 1.3); while in other areas or for minority people, the birth quota is two, which is shown in Figure 1.4.

Although the two-child policy results in a slightly complicated decision process, the structure of the process can be summarized as a four-step decision tree. First, parents decide whether to have children; second, for those who find the first kid is going to be a girl, they need to choose whether to engage in sex selection; third, parents decide whether they will have a second child; fourth, those who decide to have a second child and realize the second one is going to be a girl need to decide if they will abort this girl.

Before proceeding, we would like to discuss one subtle question: whether the strategy of “abort-and-reconceive-until-a-boy” could be applied for the second time. As seen from Figure 1.4, in one decision route parents abort the first girl until having a boy, then decide to have a second kid and find it is going to be a girl again. At this moment, these parents may have the option to take abortions again until a second boy is coming. However, we eliminate this option for them because (1) nobody can have an unlimited number of conceptions and abortions in their lifetime, (2) the price (both fixed and time cost)

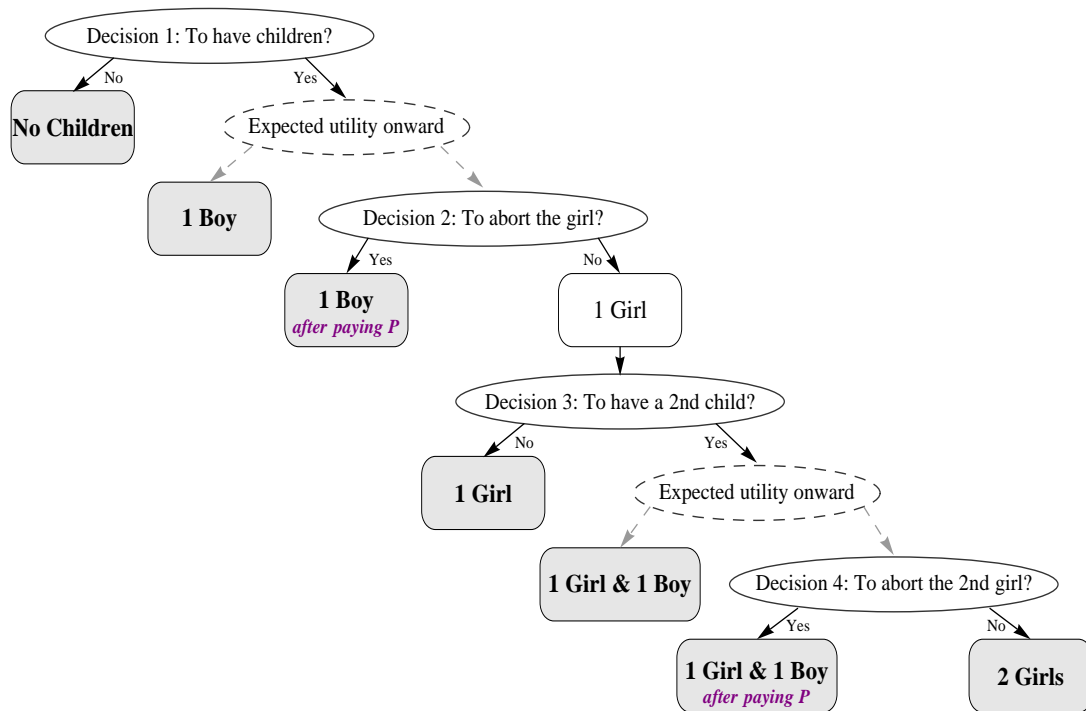


Figure 1.3: Decision Tree under a One-and-Half-Child Policy

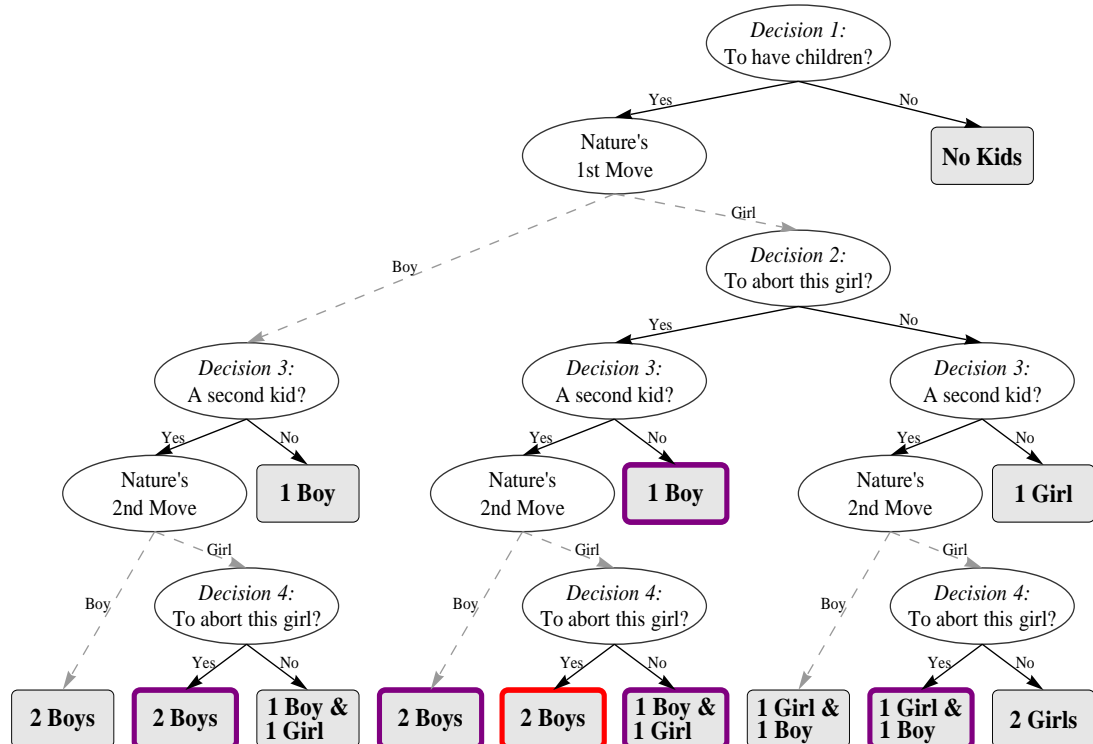


Figure 1.4: Decision Tree under a Two-Child Policy

of applying the sex selection technology for the second time may be much higher than for the first time, and (3) after trying several times to get the first boy, the probability of successful “abort-and-reconceive” will decrease dramatically. Therefore, we impose the minor restriction that for those who arrive at this decision node, it is a one-shot decision that they have to accept whatever the nature’s choice is.

Now we need to solve the individual optimization problem for five additional scenarios: (1) two boys with sex selection; (2) one boy and one girl with sex selection; (3) two boys; (4) one boy and one girl; and (5) two girls. First, let’s look at the representative scenario of “two boys with sex selection, social insurance and adaptive parents” as follows:

$$\text{Max}_{\{s_t, d_t\}} \mathbb{U}(C_t^m, C_{t+1}^o, C_t^o) = \{u(C_t^m) + \delta \mathbb{E}_t u(C_{t+1}^o) + \eta u(C_t^o)\} \quad (1.13)$$

s.t.

$$C_t^m = (1 - s_t - d_t)(1 - \alpha_t)W_t - 2(a + b(1 - \alpha_t)W_t) - (c + \phi W_t) \quad (1.14)$$

$$C_{t+1}^o = \left(\sum_{i=1}^2 d_{t+1,i} w_{t+1,i} \right) (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \mathbf{SI}_t \quad (1.15)$$

$$C_t^o = f_{t-1} d_t (1 - \alpha_t) W_t + R_t s_{t-1} (1 - \alpha_{t-1}) W_{t-1} + \mathbf{SI}_{t-1} \quad (1.16)$$

$$\mathbf{SI}_t = (\gamma_t \bar{W}_t + \beta_t \alpha_t (1 - 2b) W_t) R_{t+1} \quad (1.17)$$

$$\mathbf{SI}_{t-1} = (\gamma_{t-1} \bar{W}_{t-1} + \beta_{t-1} \alpha_{t-1} (1 - f_{t-1} b) W_{t-1}) R_t \quad (1.18)$$

The solution procedure for this model is quite similar to the one-child case and we present the details in Appendix 1.7.2. Now, let’s look at other variants which are slightly simpler. Basically, the version without social insurance will be that $(\alpha, \beta, \gamma) = (0, 0, 0)$;

the one without sex selection will be that $(c, \phi) = (0, 0)$; the one involving girls, like “one boy and one girl” and “two girls”, will be that $(\alpha, \beta) = (\alpha_g, \beta_g)$ and $(\pi_g, \lambda_g) = (0, 0)$.

1.4 Model Results

1.4.1 Parameter Choices

Our benchmark calibration aims to show the direction and magnitude of the three factors in the benchmark year. Throughout the calibration, we assume that a period is 20 years which may be a bit away from the realistic “40-years-working and 20-years-retirement,” but is helpful to choose the benchmark year. As seen in Figure 1.1, the sex ratio increased dramatically during the period 1980-2000, while it seemed to stabilize at a high level after 2000. On the one hand, China’s social insurance programs were launched around 1993. The reforms of these programs in urban areas mostly took place in early 2000s; expanding it to rural areas is a more recent development. On the other hand, the family control policy was not a heated topic in 2005, but policy makers and researchers intensively debated the pros and cons of relaxing the policy starting in 2009. For both reasons, year 2005 might be regarded as a year without dramatic social insurance and family planning policy changes. Details on the benchmark parameters are presented in Appendix 1.7.3. Here we provide formal justifications for those benchmark parameters related to child investment.

First, childbearing cost consists of a fixed cost (a) and an income-varying cost (bW_t). We assume a is around 5% of average income, i.e. one-year’s average income, and set b as 10%. Providing support for our assumptions, Echevarria and Merlo (1999)

find that the cost to a woman of having a child is about 5% of her working lifetime. Juster and Stafford (1991) find that hours per week allocated on childcare account for between 6.43% and 18% of parents total available time. In China, children are heavily dependent on their parents' support. In addition, we make no distinction between boys and girls on childbearing cost, and we assume that raising two children will double the cost, without considering any economies of scales.

Second, sex selection cost consists of a fixed cost (c) and an income-varying cost (ϕW_t). This price reflects the expectation of the accumulative cost of having a series of girl abortions until a boy is coming. Costs not directly related to one's income such as the screening test and surgery fees are captured in the fixed part. Since several abortions might be needed in order to conceive a boy, we assume that the fixed cost equals one-year's average income. However, we simplify the income-varying cost. (1) For each conception, couples need to wait for at least 4 months to know the gender of the baby. (2) If a couple decides to abort, the woman needs time to recover before they can have another try: we assume that the recovering time is 3 months. This is a moderate assumption and we admit that some only wait for 1 or 2 months, while others wait much longer. (3) The probability of consecutively having girls is decreasing as a geometric series: the chance that the first try is a girl equals $\frac{1}{2}$, the first two are girls equals $\frac{1}{4}$, the first three are girls equals $\frac{1}{8}$, and so on. (4) A woman can have at most 4 abortions in her life. Given that time cost for one abortion is 7 months (4 months on waiting for screening and 3 months for recovery), the expected time cost is 11.375 months ($= \frac{1}{2} \times 7 + \frac{1}{4} \times 14 + \frac{1}{8} \times 21 + \frac{1}{16} \times 28$). Roughly speaking, a couple expects to commit 12 months to have a boy, hence we assume $\phi = 0.05$.

Third, we assume that children’s transfer rate follows a Beta distribution and relative income is log-normally distributed. The distribution of transfer rate from boys and girls will reflect the “high risk high return, low risk low return” property. This is intuitive because the economic aspect of son preference is that sons can provide more financial support to elderly parents, while daughters, after getting married, normally will not live with their parents and hence provide less support. However, transfer from sons might be affected more by other factors and be more volatile. Therefore, we assume $\mathbb{E}_t[d_{t+1,\text{boy}}] > \mathbb{E}_t[d_{t+1,\text{girl}}]$ and $\text{Var}(d_{t+1,\text{boy}}) > \text{Var}(d_{t+1,\text{girl}})$ ²⁰.

Last but not least, relevant for families who take the “can-not-marry” risk into consideration, the transfer rate from an unmarried son is still Beta distributed, but with a mean discounted by 25% of that of his married peers. Such discounting could be justified that if a male can not marry, he will not have children to support himself in his own retirement. Rationally this single male needs to save more for himself and transfer less to his parents, other things equal. Correspondingly, the variance of single male’s transfer rate equals $0.75^2 \text{Var}(d_{t+1,\text{married boy}})$ so that the coefficient of variation is the same between married and unmarried sons.

1.4.2 Benchmark Calibration

In our benchmark models, parents are characterized by three dimensions of heterogeneity: their own income $W_{t,i}$, their parents’ (i.e. grandparents in our model) income

²⁰Here we assign values for the transfer parameters in an exogenous way. This can be improved in two ways. One is to do sensitivity analysis on these parameterizations. The other is to endogenize the transfer distribution so that the prior expected transfer distribution coincides with the actual distribution of optimal transfers from parents. We pursue the second route in Section 1.5.2.

$W_{t-1,j}$, and the expected income distributions of their kids $W_{t+1,k}$. First, for every (i, j, k) pair, we solve the individual maximization model for each terminal scenario, like the “two boy with sex selection, social insurance and myopic parents”. Second, for each combination of (i, j, k) , by comparing the maximal utility at each decision node, we obtain the optimal fertility, saving, and transfer choices. Third, we aggregate the individual fertility choices over dimension k , then over j and finally over i , which leads to a society-wide sex ratio, corresponding to a particular setting like “one-child policy with social insurance and myopic parents”.

While there is no confusion on weighting schemes for i and j (the empirical distribution of own and parents’ income), the weighting scheme for k is a bit complicated: basically we employ a Markov transition matrix from the parents’ income position to the (expected) children’s income position. To simplify, we assume that the grid points in children’s income distribution are the same as the parents’, and parents could attach different probability combinations to these grid points. We consider two possible ways to construct the transition matrix. The first one is that all parents have an identical expectation on the income distribution of their children, which is the same as that for the current parents’ generation (denoted as the *Raw* matrix). The second one is that parents’ expectations on children’s income are correlated with their own income position (denoted as *Adj* matrix). Intuitively, a millionaire should use a weighting scheme that put a higher probability on the high income percentile rather than the probability from the empirical society-wide income distribution. We conjecture that the *Raw* weighting matrix will provide a lower bound on the simulated sex ratio, while the *Adj* version will provide an upper bound.

Table 1.2 reports the calibrated sex ratio in the society, in which columns labeled

Table 1.2: Calibrated Society-Wide Sex Ratio

		Myopic		Adaptive	
		<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>
SI-No	One-Child	113.7	135.6	104.0	115.3
	One-N-Half-Child	102.5	108.8	100.8	103.9
	Two-Child	103.1	112.9	100.6	102.9
SI-Yes	One-Child	119.4	147.7	105.2	117.8
	One-N-Half-Child	102.6	109.0	100.8	104.0
	Two-Child	104.4	115.5	101.9	109.2

as *Raw* and *Adj* correspond to the two different weighting schemes on expected kids' income. Several important observations stand out.

First, relaxing the one-child policy will significantly alleviate the sex imbalance, which is consistent regardless of the presence of a social insurance program and parental expectations. Basically, giving some or all parents a second chance will ease their intention of distorting the sex ratio of the first birth; most of those who choose sex selection under the one-child policy, will now wait until the second child to engage in sex selection. As discussed above, the values using *Raw* and *Adj* weighting schemes may correspond to the lower and upper bound of the simulated society-wide sex ratio, their average under the one-child policy for myopic parents without social insurance is around 125, which decreases to 106 under the one-and-half-child policy. On the comparison between one-and-half-child and two-child policy, the difference in sex ratio is not dramatic. The underlying explanation is quite subtle, and the key is to understand the behavior of those parents, who already have a son and now are allowed to have a second child. On the one hand, some of them may choose to do sex selection on the second birth; on the other hand, they should have a smaller incentive to distort the sex ratio of the second birth than those parents who

already have a daughter²¹.

Second, the impact of social insurance on the sex ratio is a bit surprising: when social insurance is present, the sex ratio increases under every policy environment. The magnitude of the increase is limited under the one-and-half- and two-child policy, but is moderate in the one-child case. Due to the price and income effect, social insurance taxes make child care cheaper, and pension benefits increase the life-time income for the majority of people. However, we also acknowledge that there are two additional unconsidered channels. The first one is a liquidity constraint. Social insurance might decrease a couple's current cash-on-hand, even if it raises their life-time income, under the implicit assumption that the financial market is incomplete and people can not borrow against their pension benefits. The second is that we assume parents haven't taken into account the possible social attitude change that their kids may not transfer as much as otherwise. We suspect that these two channels have worked at this moment, given that social insurance is still underdeveloped, but we will explore these two issues explicitly in section 1.5.3.

Third, when parental expectation shifts from myopic to adaptive, it also helps reduce the sex ratio. For example, under the one-child policy without social insurance, the average of the simulated sex ratios using *Raw* and *Adj* weighting schemes falls from 125 for myopic parents to 110 for adaptive parents. The direction of this effect is consistent

²¹Suppose among the first birth, the total number of kids is w_1 and the percentage of boys is z_1 ; and the corresponding variables are (w_g, z_g) among the second birth whose first birth is a girl, and (w_b, z_b) whose first birth is a boy. Then the percentage of boys among all newborns under the one-and-half-child policy is $(w_1 z_1 + w_g z_g)/(w_1 + w_g)$, and that under the two-child policy is $(w_1 z_1 + w_g z_g + w_b z_b)/(w_1 + w_g + w_b)$. One can show that whether two-child policy is better than one-and-half-child version depends on whether z_b is smaller than $(w_1 z_1 + w_g z_g)/(w_1 + w_g)$. Since the constraint we have is $100 \simeq z_1 \leq z_b \leq z_g$, which is weaker than the above condition, it is hard for us to predict which policy is better.

with our conjecture and the magnitude is significant across all scenarios. When parents take into account the “can-not-marry” risk and the correspondingly lower transfer from unmarried sons, as an investment vehicle sons are not as attractive as otherwise and parents will adjust their fertility choices. It might be the case that some couples at the margin of choosing sex selection, would choose not to invest in such a less rewarding asset. Thus we see fewer distortions and a better sex ratio.

Last, as expected the *Adj* weighting matrix generate a higher aggregated society-wide sex ratio. This fact indicates that, other things equal, high income people, with prospects of having rich children, are more likely to do sex selection and distort the gender balance. Thus when heavier weights are applied on such prospects, we will see a more unbalanced sex ratio.

1.4.3 Heterogeneity with respect to Parents’ Income

The benchmark calibration gives us an overall idea on how the three factors could possibly affect parents’ fertility choices. However, it is natural that their impacts could differ from person to person. In this section, we will focus on the effect of parents’ income on fertility choice by looking at sex ratios for each income subgroup. Everything here is similar to the benchmark case, expect that the sex ratio within a given level of parents’ income is a statistic by taking the aggregation over the dimensions of grandparents’ ($W_{t-1,j}$) and expectation on children’s income ($W_{t+1,k}$).

1.4.3.1 The Relationship between Sex Ratio and Parents' Income

Figure 1.10 shows how the sex ratio changes with respect to parents' income. However, at first sight, the result seems unintuitive that the figure lacks a consistent pattern: there is neither a monotonic nor an inverted U-shaped relationship. For example, one could spot several sharp up-and-downs from the red line in Figure 1.10(a): the sex ratio first increases from 100 to 300 when parents' income rises from $1.1\bar{W}_t$ to $2.5\bar{W}_t$, then declines to 160 for those who earn $[2.6, 3.3]\bar{W}_t$, but later goes back to 300 when the income reaches $3.4\bar{W}_t$, and further becomes infinity (i.e. all boys and no girls) for parents earning $[4.9, 5.7]\bar{W}_t$. These jumps cast doubt on our model and solution. Below using the one-child policy as an example, we explain these seemingly strange results and show that there is in fact a very good intuition behind it: namely, a declining TFR and an inverted U-shaped sex selection decision are jointly responsible.

First, let's recall that in the decision tree under one-child policy (Figure 1.2), any distortions in the sex ratio come from the possibility that some parents decide to abort the incoming girl. And their optimal decision depends on the comparison of maximal life-time utility from "keep the girl" (i.e. "1Girl") and from "abort the girl until a boy is coming" (i.e. "1BoySex").

Figure 1.5 shows for a *given* combination of $W_{t-1,j}$ and $W_{t+1,k}$, the maximal utility when parents keep the incoming girl (the dashed line) and when they pay sex selection cost to have a boy (the plain line)²². Two points can be made here. First, although the

²²For illustrative purposes, we choose $j = 8$ and $k = 8$. The solution of the model shows that, with an expectation of higher children's income, given childbearing cost and children's transfer rate, the gap in the rate of return between boy and girl investment enlarges and boys become more attractive to parents. If we instead choose a smaller k (i.e. children's income are distributed around a lower percentile), we will see that the plain line is always below the dashed, which means nobody will abort the girl given that low-income




Figure 1.5: Maximal Utility: 1BoySex v.s. 1Girl

fertility choices are discrete and we solve the model for a grid of $W_{t,i}$, the (linearly) interpolated maximal utility for either decision is well behaved: it is increasing, concave and continuous in parents' income. This suggests that the numerical solution of the model is correct. Second, since the plain and the dashed lines are very close in Figure 1.5, we look at their difference in Figure 1.6, which is also continuous. When parents' earning is between $W^{\#1}$ and $W^{\#2}$, aborting the incoming girl is a better choice than keeping her. This is again in line with the intuition that only those who are neither too poor nor too rich can afford both the fixed and income-varying costs of sex selection.

Second, although imbalances in the sex ratio come solely from parents' decisions at the second decision node, one must also look at how many people could possibly arrive at that node, because some parents may simply decide not to have children. Figure 1.7 shows the maximal utility between having and not having children, and is of similar shape to Figure 1.5. In terms of making optimal decisions at the first decision node, the key is to compare the plain and the dashed lines. In Figure 1.7(a), *given* an expectation of a higher boys are not that attractive.





Figure 1.6: Maximal Utility Difference Between 1BoySex and 1Girl Decisions

children's income, the maximal utility of having a child is always larger than that of not having, although the gap diminishes as income rises. While in Figure 1.7(b), if parents are less optimistic, it is optimal to have children only for those with earnings less than $W^\#$. The existence of a threshold level of income in this example, is consistent with the declining TFR as income rises.



(a) $j = 8, k = 8$

(b) $j = 8, k = 5$

Figure 1.7: Maximal Utility: Whether to Have Children

Third, the simulated sex ratio by parents' income is an aggregated ratio over two dimensions: income of grandparents ($W_{t-1,j}$) and of children ($W_{t+1,k}$). Its value is af-

affected by how many parents decide to have a child and how many of them choose to do sex selection when facing an incoming girl. In Figure 1.8, we show a disproportionate sex selection pattern, which is the driving force for the up-and-downs of the sex ratio.

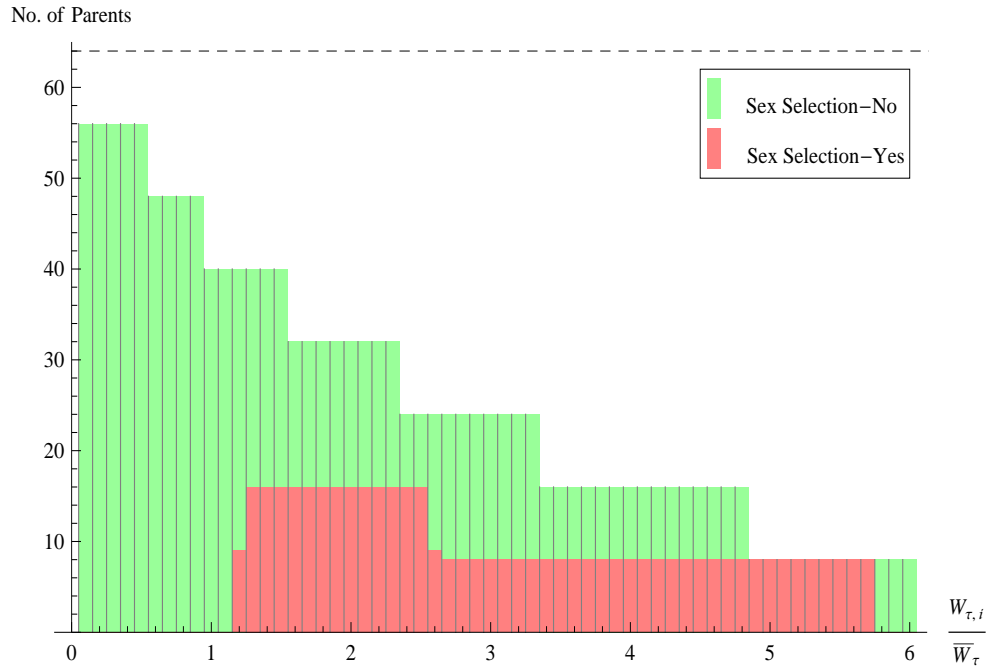


Figure 1.8: Disproportionate Sex Selection Pattern

We assume that for parents at each income level, there are 8 possible incomes for their elderly parents, and 8 possible expectations on the income of their children, thus we denote that there are 64 types of parents at a given income level, differing in their $W_{t-1,j}$ and expectation of $W_{t+1,k}$. So, the dashed gray line in Figure 1.8 is a horizontal line at 64, i.e. the maximal possible number of children within each income group. The height of each rectangle represents how many of the 64 types of parents at a given income position decide to have a child. There are two parts in the rectangle: arriving at the second decision node, the height of the green part indicates the number of parent-types who will not abort the girl, and that of the red shows the number of those who will. On the one hand, with

the rise in $W_{t,i}$, fewer types of parents choose to have a child, i.e. a declining TFR. On the other hand, the relationship between the number of parents choosing sex selection and income is inverted U-shaped, as is similar to that in Figure 1.6.

More importantly, Figure 1.8 serves as an intermediate step to compute the sex ratio. Suppose among the total 64 types of parents for each income group, Z types of parents decide to have a child. If all of them are facing an incoming girl, X types of parents (the height of the red part) will do sex selection and the remaining $Z - X$ (the height of the green part) will not. Given the nature's choice, the total number of boys born within this income group equals²³ $\frac{Z}{2} + \frac{Z}{2} \times \frac{X}{Z} = \frac{Z}{2} + \frac{X}{2}$; the total number of girls born is $\frac{Z}{2} \times \frac{Z-X}{Z} = \frac{Z-X}{2}$. Thus sex ratio is $\frac{\frac{Z}{2} + \frac{X}{2}}{\frac{Z-X}{2}} = 1 + \frac{2}{\frac{Z}{X} - 1}$. It is clear that both the number of parents who choose to have a child and that who choose to abort the girl affect the sex ratio; in fact, sex ratio is increasing in X/Z .

A more intuitive story is in place. We treat the set consisting of 64 types of parents at a give income level as a bottle. There are three materials inside the bottle: *air* which refers to those types of parents who decide not to have a child; *soil* which represents those who decide to have a child and will abort the girl if nature leads them to the second decision node, and *water* which denotes those who choose to keep the girl. Sex ratio then can be interpreted as the soil concentration in the *solution*. First, *air* does not affect the cleanness of the solution, just like those not having children contribute nothing to the unbalanced sex ratio. Second, the size of the *solution* approximates the total number of parents who decide to have a child, i.e. Z . Third, *soil* directly affects how muddy the

²³For illustrative purposes, the following is an unweighted simplification. In our baseline calibration, we also weight the contribution of each j and k by their proportions in the income distribution.

solution is, and this closely resembles the following: sex selection behavior determines to what extent the sex ratio is distorted. Since the soil concentration could be computed as X/Z , it is not surprising to realize that sex ratio is positively related to X/Z .

Therefore, changes in the size of *solution* and *soil* help explain the up-and-downs in the sex ratio. Several cases are possible. (1) Sex ratio stays constant because the sizes of both do not change. (2) If sex ratio rises, then either we have more *soil* or less *solution* or both. For example, the rise of sex ratio when $W_{t,i}$ increases from $1.2\bar{W}_t$ to $1.3\bar{W}_t$ is because more people say “Yes” to sex selection while the total size of *solution* does not change. However, the rise in sex ratio when $W_{t,i}$ moves from $2.3\bar{W}_t$ to $2.4\bar{W}_t$ is a different story: the size of *soil* does not change, but some *water* at $2.3\bar{W}_t$ now evaporates into the *air*). (3) Sex ratios become less biased thanks to a smaller size of *soil* such as when $W_{t,i}$ changes from $2.5\bar{W}_t$ to $2.6\bar{W}_t$.

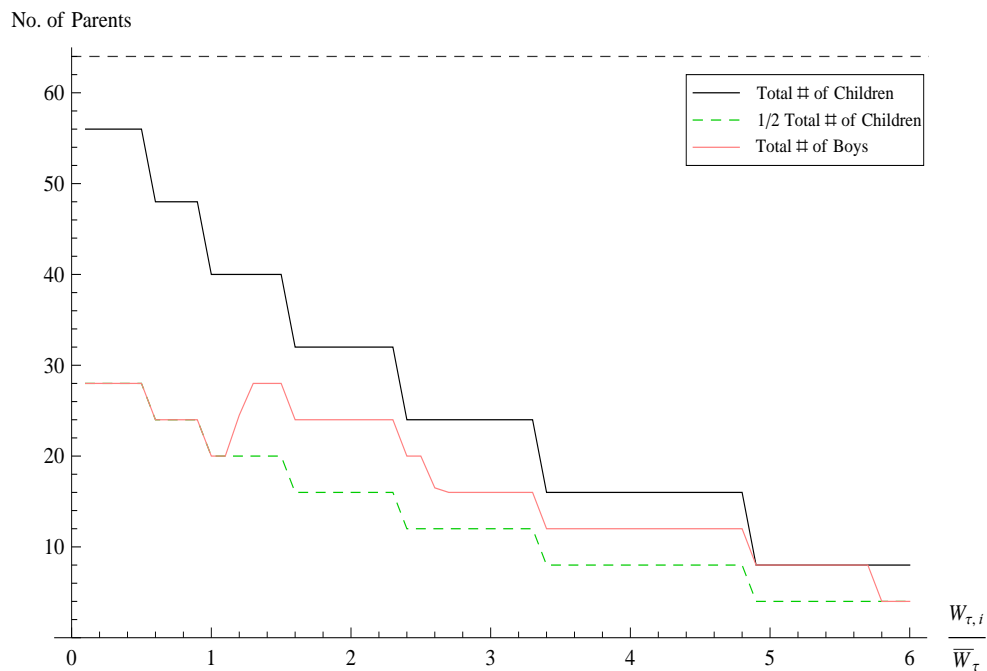


Figure 1.9: Extra Boys from Sex Selection

Finally, another angle is to look at the extra boys from sex selection, i.e. the gap between the red and the green line in Figure 1.9. Here, the black line is the total number of children born; the dashed green line indicates if nobody choose to do sex selection, the balanced number of boys within each income group; and the red line shows the actual number of boys. For some income groups, the red and the green lines coincide with each other, which means nobody say “Yes” to sex selection and the sex ratio is balanced. But for most income groups, the red line is above the green line, indicating extra boys born by aborting incoming girls. It is obvious that the magnitude of the differences between these two lines depends on the position of $W_{t,i}$. We can then calculate sex ratio directly by dividing the height of the red line by the gap between the red and the black lines.

To sum up, our two fundamental intuitions are correct that (1) TFR is a decreasing function of parents’ income, and (2) in face of an incoming girl, the number of parents engaging in sex selection will first increase and then decrease as income rises. In other words, we have Z as a decreasing function and X as an inverted-U function of $W_{t,i}$, so it is natural that $\frac{X}{Z}$ would be neither a decreasing nor an inverted-U function of $W_{t,i}$. Since sex ratio is an increasing function of $\frac{X}{Z}$, we in turn see a not-so-well-behaved pattern of sex ratio as income rises.

1.4.3.2 Heterogeneous Impacts of the Three Factors

We now compare the sex ratio by income group under various scenarios and the results are shown in Figure 1.10. First, although the overall magnitude of relaxing the one-child policy is significant, people are not equally affected. Under the scenario that

parents are myopic and not covered by the social insurance program (as shown in Figure 1.10(a)), those at the lower ($\leq 1.1\bar{W}_t$) and higher ($\geq 5.8\bar{W}_t$) end of income distribution have no response to the hypothetical policy changes, because both types of people will not choose sex selection and thus the sex ratio of children from them is always 100. However, myopic parents with income between 1.2 and 5.7 times the average level, who can afford both the fixed and the income-varying parts of the sex selection cost, are affected most by the policy change. Even within this income range, the impacts of relaxing the policy are heterogeneous: the magnitude becomes larger when people earns more. For example, for parents earning 4.8 to 5.7 times the average, the sex ratio is simply infinite under one-child policy, which indicates that they will continue to “abort-and-reconceive” until a boy is coming, regardless of any cost considerations, while it declines dramatically to 200 under the one-and-half-child policy. For parents earning 1.6 to 1.8 times the average, the sex ratio drops from 200 to 118 if they could have a second child if the first is a girl, which is a relatively small amount *compared to the previous example*.

Second, the limited effect of social insurance on the sex ratio is confirmed by this heterogeneity analysis. Figure 1.10(b) is an example to illustrate this under the one-and-half-child policy for myopic parents. It is quite clear that only people at particular borderline incomes are affected. Basically social insurance brings two effects: a smaller childbearing cost for everyone (price effect), and a higher life-time income for the poor but a lower one for the rich (income effect). Although we do not assume a tax-deduction in sex selection cost, the enforcing price and income effects for the poor imply that their childbearing and sex selection motives will be strengthened, and whether sex ratio will rise or decline depends on which of the two motives grows faster. For the very rich, the

case is that the income effect dominates the price effect, so sex ratio changes for them are determined by relative reductions in both motives.

Third, a shift in parental expectation affects a much broader range of people: as long as parents distort their fertility choices, taking into account the “can-not-marry” risk could help correct their distortions to various degrees. For instance, for myopic parents under the two-child policy environment (as shown in Figure 1.10(c)), almost all parents within the income range $[1.5, 5.7]\bar{W}_t$ change their decisions on sex selection if they realize the “can-not-marry” risk. The most striking changes are for couples with earnings between $4.9\bar{W}_t$ and $5.5\bar{W}_t$, who without recognizing the “can-not-marry” risk for sons, have huge incentive to abort, causing an extremely abnormal sex ratio (700). However, these people will not distort their fertility choices at all if they expect a difficult time for sons to find wives.

1.4.4 Sensitivity with respect to Sex Selection Price

Our structural model is designed to explain the fertility choices under the assumption that children are a type of investment vehicle, thus those parameters measuring the investment cost and return play a key role in explaining our qualitative and quantitative results. Investment cost involves sex selection price ($c + \phi W_t$) and childbearing price ($a + bW_t$); meanwhile the return side is reflected as the expected transfer in the future. Here we focus on the sex selection price, which affects peoples’ choices on whether to “abort-and-reconceive,” and hence the sex ratio directly²⁴.

²⁴We also acknowledge that the childbearing cost will directly affect the fertility number (TFR), but not sex selection decisions.

1.4.4.1 Fixed Cost

The fixed cost of childbearing reflects the total cash needed in a series of abortions and re-conceptions until a baby boy is coming; and the main component may be the ultrasound detection and surgery expenses. We set the benchmark as $0.05\bar{W}_t$ (one year's average income), and then double and half the cost (i.e. 6 months' or 2 years' average income). The simulated sex ratios are reported in Table 1.8 and the cross-sectional differences in the response are represented in Figure 1.11.

Several observations are worth mentioning. First, the sex ratio is strongly negatively correlated with the fixed cost. When c rises from 6 months' average income to 24 months', the sex ratio drops dramatically, especially for myopic parents. Second, peoples' responses to the change in c are heterogeneous. Some dramatically adjust their choices: for example, under the one-and-half-child policy without social insurance, myopic parents earning $5\bar{W}_t$ end up with an abnormally high sex ratio (200) when c equals 6 months' average income; while the same couples will have a normal sex ratio when c rises to $2\bar{W}_t$. Third, when the fixed cost equals 2 years' average income, almost everyone in the society will consider this price tag to be extremely expensive and not intervene the nature's choices. Suppose the annual average income is \$50,000, then "24 months" means that the expected medical cost, associated with repeated ultrasound detections and abortions, is \$100,000, which is indeed prohibitive, and perhaps even the richest may not adopt such strategy.

The third observation has particular implications for policy makers in an effort to re-balance the sex structure. Relaxing the one-child policy could alleviate the imbalance,

but the side-effect is that TFR will increase simultaneously. In the meanwhile, we see that social insurance's influence is quite limited. In light of these observations, one solution is to raise the fixed cost for abortions. Although the official regulation claims that non-medical-related abortion is illegal, the enforcement is weak and sex-selective abortions are quite common with low ultrasound and surgery charges, especially in less developed areas. If the government could strengthen their monitoring and impose fines for violation (i.e. a way to increase the fixed cost of sex selection), China would see a re-balancing process of its sex composition. To sum up, a penalty along with a strict enforcement could stop most people from choosing sex selection. However, we also admit that this is equivalent to using exogenous forces to alter peoples' budget constraint and it is always easier said than done. A more desirable policy should influence individual's preference or expectation, so that couples' optimal decisions coincide with social planner's objectives, without much intensive monitoring from the government side.

1.4.4.2 Income-Varying Cost

This part reflects the accumulated time involved in a series of abortions and re-conceptions and we set the benchmark as 12 months ($\phi = 0.05$), with the justification in section 1.4.1. The simulated society-wide sex ratios are reported in Table 1.9. Since ϕW_t affects the child investment return in a negative way, we also observe a significantly negative correlation between the income-varying cost and sex ratio. For myopic parents, when ϕ increases from 12 months to 15 months, we observe a close-to-normal sex ratio in every case except the one-child case using *Adj* weighting matrix, which indicates a rather

powerful channel.

We also notice the difference in cross-sectional distribution of such effect, which is displayed in Figure 1.12: when the income-varying cost increases, the sex ratio drops more for high-income couples, but remains unchanged for low-income ones. The underlying reason is straightforward: since this cost is proportional to one's income, the rich will feel the price pressure more. For example, parents earning $5\bar{W}_t$, with c moving from 9 months to 15 months, see their fixed part expense as a share of their income increases from 0.75% to 1.25%; and when ϕ rises from 9 months to 15 months, their income-varying expense jumps from 3.75% to 6.25%, a larger price hike. However, for parents earning $0.5\bar{W}_t$, with the same change in c and ϕ (9 months to 15 months), the fixed part will rise from 7.5% to 12.5%, which is larger in both absolute and relative terms; while the change in income-varying part remains the same (3.75% to 6.25%).

As a summary, the sensitivity exercise shows that: the calibrated sex ratio are sensitive to changes in sex selection cost, especially in the one-child case; and across the section, the poor are affected more by the fixed part, while the rich by the income-varying cost instead.

1.4.5 Discussion on Model Fitting

Based on analysis so far, we know that there are various dimensions of heterogeneity such as policy environment, parents' income, sex selection cost and so on, that can have impacts on the aggregated sex ratio. Given that all these dimensions can change simultaneously in real life, to see if our theoretical model could produce a close match

with actual sex ratios, we need to make sure variations in all these dimensions are taken into account in a sensible way.

We first describe how the empirical statistics look like, then discuss potential contributing factors, and demonstrate the correspondence between the actual and simulated fertility outcomes.

1.4.5.1 Data

Our population census samples span three census years (1982, 1990, and 2000), which provide a unique opportunity to assess the time-series and cross-sectional responses of fertility outcomes to changes along several dimensions including the family-control policy, social insurance coverage, sex selection and childbearing costs, and so on. The three columns labeled by “Actual” in Table 1.3 reports the actual sex ratios of different socioeconomic groups from the three census years²⁵.

We use individual’s response on ethnicity to identify whether a person is an ethnic minority. Although China has 56 official recognized ethnic groups, the family planning policy does not differ much among those who are not Han Chinese, and thus we group them as “minorities”. Meanwhile, due to data limitation, the rural and urban classification is less clear. The regulation stipulates that whether a couple can have a second child depends on their *hukou*²⁶ status: people with agriculture *hukou* are under the one-and-half-child or two-child policy, and people with non-agricultural *hukou* are subject to the

²⁵There is a slight discrepancy in 2000 data between this table and Table 1.1, where the summary statistics are based on full census sample.

²⁶In China, every household is required to register in the residence registration system, which was officially promulgated by the government in the 1950s to control the movement of people between urban and rural areas. A *hukou* (i.e. a household registration record) identifies a person as a resident of an area and includes identifying information such as name, parents, spouse, and date of birth

one-child limit. So we use the *hukou* status to identify rural and urban residents²⁷.

Table 1.3: Actual versus Simulated Sex Ratio

	1982		1990		2000	
	Actual	Simulated	Actual	Simulated	Actual	Simulated
Overall	107.9	108.7	113.0	113.0	120.4	123.0
Han	108.4	108.9	113.7	113.6	121.8	123.7
Minorities	102.8	105.0	107.9	106.1	111.6	115.0
Rural	107.8	108.9	113.0	112.6	125.0	125.7
Rural Han	108.4	109.1	113.7	113.2	126.4	126.5
Rural Minorities	103.0	105.0	107.2	106.4	117.1	117.3
Urban	108.1	108.0	113.6	113.9	117.1	118.3
Urban Han	108.3	108.2	113.4	114.7	118.8	119.0
Urban Minorities	101.2	105.0	116.8 [‡]	105.0	100.0 [‡]	110.9

Note: These two statistics may be subject to small sample errors. A reasonable estimate for the Urban Minorities should be 107 in 1990 and 110 in 2000.

Source: China population census in 1982 (1% sample), 1990 (1%), and 2000 (0.1%).

Several patterns of actual fertility outcomes are worth noting. First, there is a clearly increasing tendency of sex ratio from 1982 to 2000. The overall sex ratio rises from the close to normal level (107.9) in 1982 to an abnormally high level (120.4) in 2000. This trend is consistent across all subgroups, with the biggest increase (16.6% increase) among Han Chinese living in rural areas. Second, sex ratios for Han Chinese are significantly higher than that for minorities and the gap between the two has widened as time goes on. This observation is consistent with the fact that minorities were exempt from the family control regulations, but more Han Chinese faced strict fertility limits from 1982 to

²⁷ *Hukou* information is not available in our 1982 sample, so we use a person's living place (whether in a city or in a prefecture) as a proxy. Furthermore, one minor issue with the 1990 sample is that we find around 20% newborns with missing *hukou* status. This might be due to a temporal revenue-generating policy around that period, which allowed rural residents to pay "fees" to transfer their (and their children's) *hukou* type. Given the advantages of holding a non-agriculture *hukou* on housing, education, employment, health care, and so on, relatively rich rural people had great incentives to make such a transition. However, the whole process took a long time. Those newborns without a clear *hukou* type were most likely in such a transition process, and hence we classify them as rural residents.

2000. Third, although the sex ratios in rural and urban areas are quite close in both 1982 and 1990, it became more unbalanced in rural areas in 2000, implying some underlying changes having taken place during this period. Finally, the higher sex ratio for Han Chinese than for minorities is consistent in both rural and urban areas, and the similarity and discrepancy in the sex ratio between rural and urban areas remain the same across both ethnic groups.

There is a need for some explanation on sex ratios of minorities living in urban areas. Calculated from our census samples, the sex ratios for urban minorities were not that reasonable: 116.8 in 1990 but 100 in 2000. This might be due to the small sample size. Given the similarity in sex ratio between rural and urban areas in our 1990 census sample, we think the actual sex ratio should be around 107 for urban minorities in 1990. For the 2000 urban minorities, we take the average of that in cities (108) and in towns (112) as in Table 1.1, since cities and towns are more close to urban areas. These adjustments may also imply that compared to the full census sample, values reported in Table 1.3 were slightly higher in 1990 and lower in 2000.

1.4.5.2 Contributing Factors in the Cross Section and Over Time

Variations in a number of factors contribute significantly to the pattern of sex ratio as we see in Table 1.3. And we classify them into two types. The first type includes factors that only change their values over time but not across different population groups. One example is the fixed cost for childbearing, which rises dramatically from the early 1980s to 2000. China nowadays has a below-replacement TFR of 1.8, which can not be

contributed by the birth quota. Instead, it is due to a rising proportion of urban couples who postpone their childbearing plan and rural parents who decide to only have one child even if they are allowed to have two. A more reasonable justification is the rising child-bearing cost, which makes children too expensive to bear. To incorporate this variation, we assume that the fixed cost for raising one child equals 6 months average income in 1982, 12 months in 1990, and 18 months in 2000. Other factors that change over time include the social insurance program, which was not in place in the 1982 and 1990 census; the one-child policy, which was at its early stage of implementation in 1982, but became stricter in 1990 and 2000; and the fixed cost of sex selection, which becomes more affordable with the diffusion of ultrasound machines over time.

The second type includes factors that differ across groups. On the one hand, rural and urban areas are divided in many aspects. (1) In year 1982, the one-child policy just became effective for a few years. We assume rural residents were not affected, while only 50% of urban population were regulated by the quota. In 1990, we assume that 90% urban residents were limited by the one-child quota. And in 2000, all Han Chinese living in urban areas had to obey the regulation. Things are different for rural residents in that mothers of a daughter in several rural provinces are allowed to have a second child (i.e. the one-and-half-child policy), and families in remote areas can have two or even three children regardless of the gender of the first one. According to Ebenstein (2011), nowadays the fertility policy imposes a one-and-half-child limit on most rural residents, who account for around 54% of the total population, and a two-child limit (10% population) for remote provinces. We then assume that in 1990 and 2000, 84% (i.e. $\frac{54\%}{54\%+10\%}$) of the rural population were subject to the one-and-half-child policy, and the remaining were al-

lowed to have two children. (2) We assume social insurance coverage was available only for urban residents in 2000. (3) Rural families have historically stronger son preference because sons can provide more help on farm work and are more likely to live with their parents after getting married. Thus, we assume that the mean of the transfer rate from adult sons to their elderly parents are higher in rural areas than that in urban areas. (4) The fixed cost of sex selection, which reflects the total cash needed in a series of abortion and re-conceptions, also differs between rural and urban areas. When the ultrasound machines were not widespread, it was more expensive for rural residents to conduct ultrasound detection. However, with the diffusion of ultrasound machines, rural couples have much easier access to the sex selection technology and the cost becomes cheaper in rural than in urban areas.

On the other hand, the difference between Han Chinese and minorities is more straightforward. Minorities are excluded from the birth limit. In addition, in rural areas, the gap in old-age support between sons and daughters might be smaller for minorities than for Han Chinese because most minorities live in a small and relatively closed community and inter-ethnicity marriage is not common so that daughters after getting married still live in the same region and thus can provide more help to their parents²⁸.

1.4.5.3 Model Fitting

Our primary objective is to match the sex ratios for four basic subgroups (i.e. Han Chinese and minorities in rural and urban areas) since sex ratios for other larger groups

²⁸It is also possible that the son preference among minorities is weaker because minorities have their own culture and are not affected much by the Confucianism.

are population-weighted average of these four. For each of the four subgroups and for each census year, we first decide which factors need to be turned on and their associated magnitudes. Then, under each parameter and policy setting, we calibrate our model in the same way as described in the benchmark calibration²⁹. Since for certain years, some subgroups contain two policy environments such as 84% Han Chinese in rural areas were subjected to the one-and-half-child limit and 16% were allowed to have two in the 2000 census, we will use the weighted average of the simulated sex ratios to represent the entire subgroup. In addition, in order to calibrate sex ratio corresponding to each census year, we also update our assumptions on real interest rate, GDP growth rate, and the two state variables. After obtaining the simulated sex ratio for the four subgroups, we calculate their weighted average for the Han Chinese/minorities groups, for rural/urban groups, and for the overall population, where the weights are the actual population shares in the 2006 China Statistics Yearbook.

The three columns labeled by “Simulated” in Table 1.3 display outcomes from our model. Overall, our simulated sex ratios are very close to the actual ones: for the four basic subgroups, except for Han Chinese living in urban areas in 1990, the discrepancies between actual and simulated sex ratios are smaller than 0.8. More importantly, we also successfully match the actual fertility patterns that the sex ratio increased from 1982 to 2000, Han Chinese had a consistently higher sex ratio than minorities, and sex ratio for

²⁹As discussed in the benchmark calibration, there are two ways to aggregate over expected children’s income: one is that parents with heterogeneous income have identical expectation on children’s income (i.e. *Raw* weighting matrix); the other is that parents’ expectation on children’s income is positively correlated with their own income position (i.e. *Adj* weighting matrix). The values using *Raw* and *Adj* weighting schemes may correspond to the lower and upper bound of the simulated sex ratio and we use their average as the simulation outcome. In addition, we need to adjust the simulated result upwards by roughly 5% as discussed previously.

rural areas was similar to the one in urban areas in 1982 and 1990 census, but became much higher in the 2000 census. For example in our model, the simulated sex ratio for Han Chinese living in rural areas is 109.1 in 1982 and 126.5 in 2000, indicating a 16% increase; the simulated sex ratios for Han Chinese in both rural and urban areas are higher than that for minorities and the difference between the two widens from 3.2 to 8.1 in urban and 4.1 to 9.2 in rural areas; and both Han Chinese and minorities in rural areas have more biased sex ratios as compared to the ones in urban in 2000.

However, with respect to minorities, our simulated sex ratios do not have a close match to the actual ones, especially for the 1982 census. However, it is not appropriate to attribute the relatively big discrepancy for minorities to our theoretical model. It is known that natural sex ratio at birth is not 100, but 104-107, where the higher probability of having a boy is used to compensate for the higher infant mortality rate for males. In the 1982 census, the sex ratios for minorities in both rural and urban areas are below 104. If one believes that 104 is the lowest natural sex ratio at birth, it seems that minorities in China have a girl preference and may abort the incoming boy instead. However, a girl preference for minorities contradicts with that their sex structure was also unbalanced in the 1990 and 2000 census. Thus, we attribute the abnormally low sex ratio for minorities in the 1982 census to a small sample error. If this were true, then our model has produced a close match.

Our subsequent objective is to match the proportion of people who have another child and the proportion of males of next birth. These two proportions are used to illustrate the fitness of his model in Ebenstein (2011). Instead of solving for the maximal utility of having a boy or a girl, Ebenstein(2011) directly assumes the value a boy or a girl

can bring to the parents and later estimates the corresponding values using micro-level data. Since the payoff of a child consists of the child's value plus a random variable, it is straightforward to estimate the probability that parents engage in sex selection and the probability of having another child. Although our individual optimization framework is totally different from the probabilistic model in Ebenstein (2011), we would like to see whether our model can also perform well in this aspect.

Table 1.4: Actual versus Simulated Sex Outcomes by Sex of Existing Children

		Proportion who have another child					
		1982		1990		2000	
Parity	Sex Combination	Actual	Simulated	Actual	Simulated	Actual	Simulated
	Overall						
First	None						
Second	One boy	0.94	0.87	0.73	0.25	0.55	0.16
	One girl	0.95	0.89	0.81	0.78	0.7	0.55
		Proportion of male among next birth					
		1982		1990		2000	
Parity	Sex Combination	Actual	Simulated	Actual	Simulated	Actual	Simulated
	Overall	0.52	0.50	0.54	0.52	0.56	0.56
First	None	0.51	0.51	0.51	0.52	0.52	0.53
Second	One boy	0.51	0.50	0.51	0.51	0.49	0.54
	One girl	0.53	0.50	0.55	0.53	0.59	0.70

The actual proportions in Table 1.4 are taken from Ebenstein (2011), which are calculated for married women aged 35-40 and their matched children aged 0-18 from the 1982, 1990, and 2000 census samples. There are four fertility patterns as seen in Table 1.4. (1) Proportion of people who have another child decreases dramatically from 1982 to 2000, which coincides with the decline in the total fertility rate during the same period. (2) This proportion is bigger for parents with a girl than those with a boy, which

means families without a son are more likely to have another child, other things equal.

(3) Proportion of male among the next birth increases steadily, indicating a more biased sex ratio in 2000 than in 1982. (4) This proportion is significantly higher for parents of a daughter, especially in the 2000 census, reflecting a stronger motivation to engage in sex selection among those whose first kid is a girl.

Although our model does not directly calibrate these two proportions, we can derive them from parents' choices at each decision node. In the decision process under one- and two-child policy, people make decisions on whether to have a first (or a second) child and whether to choose sex selection for the incoming girl. Given that we have in total $I \times J \times K$ types of parents, we can calculate how many of them have another child and how many apply sex selection, and thus the proportions.

In terms of absolute values, our simulated proportions are a bit away from the actual ones. One reason might be that our framework assumes perfect enforcement of the one-child policy so that people can't pay fines to have another child, while in reality, we do observe families pay fines for violating the regulations, which may help explain the lower proportion simulated from our model³⁰. Besides this, our model matches the fertility patterns as seen from the actual proportions. The simulated proportion who have another child decreases between 1982 and 2000. Our simulated proportion of male of next birth is lower for people who already have a son than those who only have a daughter. For example, in 2000, for parents with a son, the proportion of male among the second child is 54%, while it is nearly 70% for parents whose first kid is a girl, indicating a stronger

³⁰To some extent, the proportion who have another child in Ebenstein (2011) seems to be over-estimated. By 2000, urban population, accounting for 35% of the total population, can only have one child. Even if all rural residents with a girl decide to have another child, the proportion may still be lower than 0.7.

incentive to abort the girl fetus among families with girls.

To conclude, Table 1.3 and 1.4 jointly show that our model provides a good match to the actual data; and given the model specification, our model behaves much better in terms of the sex ratios. The comparison between actual and simulated sex ratio shows that several fertility patterns and many essential elements of fertility decisions are captured by our model and confirms that our analysis on sex ratio changes when considering certain dimensions of heterogeneity is reasonable and convincing.

1.5 Three Extensions

1.5.1 Equilibrium Sex Ratio when Parents are Forward-Looking

As seen above, one robust result is that when parents realize the “can-not-marry” risk and become adaptive, the society-wide sex ratio will decline. However adaptive couples mistakenly use the sex ratio of recent newborns to approximate the sex ratio of those unborn babies. In equilibrium, forward-looking couples would calculate this risk based on expected sex ratio among the children’s cohort. The iterative procedure to find such an equilibrium is as follows: forward-looking parents will assume that every couple is rational and makes optimal decisions in the same way as they do. They first compute this risk based on a starting sex ratio κ_0 (like setting $\kappa_0 = \kappa_t$) just as adaptive couples, then make the fertility decisions. The aggregate decision from all couples behaving in this way will generate a new sex ratio κ_1 . They will compare κ_1 with κ_0 : if the two are not the same, they realize that their previous expectation either under- or over-estimates the underlying risk, and will adjust the starting point and iterate again until the realized sex

ratio (κ_{i+1}) coincides with the prior expectation (κ_i), i.e. a fixed point³¹. At this steady state, the realized “can-not-marry” risk will be identical to the expected one.

As a priori, several outcomes are possible: no equilibrium, one and only one equilibrium, and multiple equilibria. First, we expect that the no equilibrium scenario is unlikely because of the inherently monotonic property of our structural framework. For $\kappa_i > \kappa_j$, we have $\mu_i > \mu_j$ (i.e. risk under the first sex ratio is higher), the resulting sex ratio from all couples behaving rationally will have the property that $\kappa_{i+1} < \kappa_{j+1}$. That is, in an interactive setting, if the risk is exaggerated, couples will lower the starting point and if otherwise, they will adjust κ_i upwards. As iteration proceeds, the range between input κ_i and output κ_{i+1} will narrow; given a sufficient number of couples used in the numerical simulation, it is expected that the difference between consecutive κ_i will converge to 0. Second, multiple equilibria are possible depending on different starting points. However, we suspect that this is more of a theoretical curiosity rather than a realistic consideration. We will start from a neighborhood around the current sex ratio, such as $100 < \kappa_0 < 130$, and see if equilibrium exists and if multiple equilibria emerge. We define the convergence criteria as the distance between κ_i and κ_{i+1} being smaller than 0.001. Results are presented in Table 1.5.

A few observations deserve special attention. First and foremost, for most scenarios, we are able to derive one and only one equilibrium sex ratio. Second, in the benchmark calibration adaptive parents consider the sex ratio as 120 in calculating the “can-not-marry” risk and now forward-looking parents realize that the equilibrium sex

³¹A comprehensive discussion in deriving fixed point using dynamic programming method could be found at Rust (2008).

Table 1.5: Equilibrium Sex Ratio when Parents are Forward-Looking

		Forward-Looking	
		<i>Raw</i>	<i>Adj</i>
SI-No	One-Child	106.7 [‡]	117.4
	One-N-Half-Child	102.5	105.5
	Two-Child	103.1	109.0
SI-Yes	One-Child	110.6 [‡]	117.8
	One-N-Half-Child	102.6	107.5 [‡]
	Two-Child	103.2	110.4

Note: Those indexed by [‡] refer to the average value of a small range where the simulations enter into an oscillation, and convergence can not be achieved.

ratios are smaller. This tells us that adaptive parents who look at current sex ratio tend to overestimate the risk and end up in a less severe situation for their sons. As time goes by, social learning will help adaptive couples to adjust their expectations, and in the long run, the sex imbalance will automatically become less problematic. However, this learning depends on intergenerational knowledge transfer, which is not as easy and speedy as one might think. Third, we face some oscillation issues in deriving the equilibrium sex ratio for three scenarios (labeled with “[‡]”). After some iterations, the input and output κ swing back and forth between two nearby points, even if we try different starting points. Although our structural framework is monotonic, the choice variable for fertility is not continuous. Thus, it is possible that for a range of sex ratios κ , parents’ decision on sex selection remains unchanged until κ reaches a threshold value, when parents move abruptly from “no sex selection” to “applying sex selection”, or vice versa, which then induces a big jump (or fall) in the sex ratio. Under this circumstance, it is very hard to find a fixed point. In practice, we minimize the gap between the lower and upper bound-

ary for the threshold value until it is small enough, and parents are almost indifferent between doing sex selection and not. Then, we claim the middle point within this small range as the equilibrium³². Fourth, consistent with our previous findings, when parents are forward-looking, relaxing the one-child policy will significantly relieve the sex imbalance problem, and one-and-half-child policy is even better than the full-fledged two-child policy. Also, social insurance retains a positive but limited impact.

1.5.2 Endogenizing Children's Transfer Rate Distribution

In our model, middle-age people make a transfer to their elderly parents and at the same time expect their children to do the same in the future. In the above simulations, we arbitrarily pick up a specification for such expectations. One might be interested in finding some justification for this. A natural justification is: in a dynastic setting, the couple's expectation on future transfers from children is linked to the optimal transfer decisions among the couple's cohort. One example is generational independency, where parents do not care about grandparents, optimally contribute nothing to grandparents, and will not expect any transfer from their kids so that generations are independent from each other. Another example is generational generosity, where couples pay much attention to their elderly parents, and also expect their kids to care them in the same way. Both are observed in reality, perhaps due to role learning. In an infinite-horizon setting (or equivalently in a dynastic model), at equilibrium the expectation of future transfers should be identical to the cross-sectional distribution of current transfers. To endogenize the

³²Our problem is rather simple since the convergence is for one dimension only. Santos and Rust (2003) provides a more general discussion on the convergence properties of policy iteration in a class of stationary, infinite-horizon Markovian decision problems.

transfer expectation, we solve for such a steady state.

Here we match the prior expectation with an ex post realized cross-sectional distribution. In a parametric framework such as ours, this can be done only if we impose some distributional assumptions, and minimize the difference between certain parameter values. We assume that both the expectation on future transfers and the ex post distribution of realized transfers are Beta distributed, and we match the mean between the two distributions.

The reasons that we only focus on the mean convergence between two distributions come from several aspects. First, we find that there is a convergence in transfer rate for people earning above average income, indicating a small variance for the overall transfer distribution. Second, even if the convergence criteria use both mean and variance, it is hard to determine the corresponding weights for the differences in means and variances. Third, as discussed in section 1.4.1, investments on sons and daughters reflect the property of “high risk high return, low risk low return” so that $\mathbb{E}_t[d_{t+1,\text{boy}}] > \mathbb{E}_t[d_{t+1,\text{girl}}]$ and $\text{Var}_t(d_{t+1,\text{boy}}) > \text{Var}_t(d_{t+1,\text{girl}})$. However, the risk associated with child investment is not only reflected in absolute terms (i.e. bigger variance for boys’ transfer), but also in relative terms as the larger coefficient of variation for sons’ transfer. Even if both mean and variance of the transfers from sons and daughters keep changing during the iteration process, we suppose their corresponding coefficients of variation are constant, which then induces a fixed coefficient of variation for the overall child investment³³. Thus, for two

³³The argument that constant coefficients of variations (CV) for boys and girls indicate a fixed CV for the overall child investment is correct based on the assumption that the weights on boys’ and girls’ CV are also constant (i.e. the composition of boys and girls in the society does not change). In our framework, this assumption means we will not consider any composition changes due to updated sex ratio in each iteration loop.

distributions with the same coefficient of variation, if their means are close enough, their variances will also converge.

As aforementioned, we assume the transfer rate follows a Beta distribution (i.e. $d_{t+1} \sim \text{Beta}(\alpha, \beta)$) with mean $m(0)$. To incorporate son preference, we impose $m_b = 1.2m$, and $m_g = 0.7m$. Based on this, we solve our model, compute the optimal choices, and look at the cross-sectional distribution of d_t , from which we can calculate $m_{(1)}$ and then start the iteration process again. A distance measure on consecutive iterations is defined as $|m_{(i+1)} - m_{(i)}|$; we consider convergence achieved if this distance measure is smaller than a tolerance level.

As discussed previously, a couple's expectation on future transfers from children is related to their current optimal transfer decision, which in turn is affected directly by the attention they place on their elderly parents (η). One might conjecture that the equilibrium transfer is directly related to this altruism measure, which may vary from society to society. We consider this heterogeneity to show how the equilibrium is influenced directly by this altruism degree. Clearly, when $\eta = 0$, parents neither care about grandparents, nor do they expect children (if any) to care for themselves, thus the equilibrium transfer is simply zero. Table 1.10 shows the equilibrium transfer distribution for plausible altruism degrees.

First, we find that when parents become more altruistic, the equilibrium transfer distribution moves rightwards: the mean is larger without a big variation in the standard deviation. This represents the significant impact of η on intergenerational transfers. Second, the standard error is small for every scenario, indicating the overall distribution is quite concentrated. Third, only the social insurance program has a limited

effect on the distribution, while the one-child policy and parental expectation are quite silent with respect to transfer distribution. Finally, compared to our benchmark assumption ($m = 15\%$), we see that the mean of the endogenized transfer is 0-4% lower when $\eta = 0.25$, but 5-14% higher when $\eta = 0.5$.

We now proceed to calibrate the society-wide sex ratio based on the equilibrium transfer distribution and the results are shown in Table 1.11. There are several interesting findings. First, comparing Table 1.2 and 1.11, we find that when children's transfer distribution is endogenously derived, gender structure is much better (i.e. a decline in the sex ratio); and for some scenarios the sex structure will be balanced. So if parents could form their expectation on future transfers from their children in a rational way, the sex imbalance would not be as severe as in Table 1.2. Second, the three factors play their robust roles even after we endogenize children's transfer distribution: relaxing the one-child policy and shifting parental expectation could alleviate the gender imbalance, while social insurance does not seem to help solve the problem, although its impact is quite limited³⁴. Third, when η rises, although the equilibrium transfer moving rightwards, the sex ratio decreases, especially for the one-child cases.

This exercise shows the interactions between parents' degree of altruism and their fertility choices. Given children's transfer distribution, a higher η induces a close to normal sex ratio. This is intuitive: at a given transfer distribution, the return on child

³⁴Social insurance increases the sex ratio for both myopic and adaptive parents when $\eta = 0.25$, which is consistent with the observation in the benchmark calibration. When η rises to 0.5, adaptive parents do not choose sex ratio with and without social insurance, while sex ratio for myopic parents even decreases in the presence of the program. When $\eta = 0.5$, the mean of the equilibrium transfer distribution increases; given constant coefficient of variation, the variance also becomes larger. That is, the dispersion of boy's transfer distribution becomes much larger while girl's transfer is more concentrated. Thus, with a higher probability boy's transfer rate is smaller than that of girl (even if the mean for boy is still higher) so that fewer people choose to do sex selection and the society arrives at a more balanced sex ratio.

investment is fixed. When parents place more attention on their elderly parents' welfare, other things equal, they will transfer more and have fewer resources left to invest in children, making sex selection unaffordable for some couples and thus reducing the sex ratio. However, when parents are rational and endogenize their children's transfer, the effect of higher altruism degree on fertility choices are two-fold. On the one hand, keeping the transfer distribution constant, a higher altruism degree directly decreases the sex ratio. On the other hand, a higher altruism degree could indirectly lead to a more unbalanced sex ratio through the channel of transfer distribution. In detail, the higher the altruism degree, the higher the equilibrium transfer, and the higher the future transfer parents expect to obtain from their children. When the mean of the equilibrium transfer increases, the gap between sons' and daughters' transfer will enlarge³⁵. This change in the relative return from investing in boys versus girls, indicates that sons become more attractive for parents, which will induce a more unbalanced sex ratio. According to Table 1.11, the direct effect of higher η seems to dominate the indirect effect so that we observe a decline in the sex ratio.

1.5.3 A Further Look at the Role of Social Insurance

Until now, we find that social insurance tends to increase the sex ratio, although its magnitude is insignificant for most cases. This is opposite to the common belief that social insurance could help reduce the sex ratio. The intuition is that by providing generous pension benefits, social insurance substitutes part of children's role: parents do not need

³⁵For example, if the society-wide average transfer rate is 10% with boys transferring 12% and girls 8%, the difference is 4%; while if this average transfer rate increases to 20% with boys and girls transfer 24% and 16% respectively, the gap will double to 8%.

to rely as much on children as otherwise. As social insurance coverage expands from urban to rural areas, we should observe a decline in the sex ratio, other things equal.

However, it seems that our model above does not capture this feature. In our model, social insurance works through two channels: it makes childbearing cheaper (price effect), while increasing the life-time income for the majority (although it reduces that for the very rich) (income effect). Now we explore two other possible channels: liquidity constraint and social attitude changes.

First, we previously assumed that financial markets are complete so that there is no liquidity constraint, which means people can borrow against their future pension benefits. This will lead to a negative private saving for those poor people whose future pension benefits occupy a large proportion of their life-time income. On the one hand, this makes some sense in that even if the formal banking sector does not accept pension benefits as a collateral, borrowing from informal channels like among relatives is common in China and could help alleviate the liquidity constraint. The major risk of doing this is the adult mortality, which is rather small for people before retirement. On the other hand, it is possible that the liquidity constraint is binding for some people. We investigate below how this liquidity constraint affects the sex ratio.

We introduce liquidity constraint as $s_t + \theta \text{SI}_t / (R_{t+1}(1 - \alpha_t)W_t) \geq 0$, where $\theta \in [0, 1]$ is the index of financial market completeness. $\theta = 1$ means agents could borrow against the full amount of future benefits; $\theta = 0$ means they can not use any portion of the benefits as collateral. With the addition of this constraint, those who previously borrow against pension benefits, will be prohibited from doing it and have a lower utility, while others are not affected. When $\theta = 0$, the sex ratio in the third panel of Table 1.12

decreases, especially under the one-and-half- and two-child policies. This is justified by the cost of childbearing and sex selection: with one child, most couples do not need to borrow; however in the case of two kids, borrowing might become more likely, and when a liquidity constraint is in place, these people have no resources to distort nature's assignment.

Second, we previously assumed that there was no change in social attitudes with respect to future transfers from children. This is rather debatable; an alternative is to assume that after social insurance is introduced, children's transfer will be lower, i.e. the distribution of both son's and daughter's transfer will move leftwards. Couples should be able to anticipate this social attitude change twenty years from now, and adjust today's fertility decision correspondingly. This could be justified in the generational account setting: with the introduction of a "Pay-as-You-Go" system, the first retiree generation benefits most since they usually have a shorter contribution history, and this might trigger a change in the social attitude with respect to transfers from children. If this possibility is taken into account by the current parents, today's sex ratio may be different.

We assume that with changes in social attitudes, transfer from children is ϑ (a discounting factor) times the one without social insurance, and son's transfer is discounted more than daughter's. Comparing the second and the forth panel in Table 1.12, we see significant declines in sex ratio, especially for the one-child cases. That is, when parents discount sons' transfer more than daughters', their motivation of distorting fertility choices weakens because boys are not as attractive as before. And more couples would think it is not worth paying a high price (sex selection cost) on a less rewarding asset (sons). This is very similar to the "can-not-marry" risk in parental expectations, which

lowers the transfer distribution for unmarried sons but not that of daughters. Both differential treatments have proved to be the key in reducing the sex ratio, which is a result we would expect given that we focus on the “old-age support” motive for childbearing.

Finally, as discussed above, social insurance could affect fertility decisions through four channels: income effect, price effect, liquidity constraints and change in social attitudes. The overall effect when these four channels work together is shown in the last panel of Table 1.12, where we see the sex ratio is perfectly balanced for most scenarios. This should be the most comprehensive consideration of social insurance’s impact on fertility choices and serve as an upper bound for its magnitude. Under this circumstance, social insurance could play a significant role in correcting the unbalanced gender structure, but it is expected that such process will take several years if not decades.

1.6 Conclusion

In this paper, we have developed and calibrated a dynamic fertility choice model to explain the current gender imbalance in China. An excess of millions of males in the marriage market and the associated social problems are hot topics in popular media and among policy makers. While the focus of discussion has tended to be on the impact of one policy environment (either the one-child policy or social insurance) on sex imbalance, relatively little attention has been given to the reverse role of the sex imbalance on fertility choices. Our analysis shed light on some of the mechanisms underlying fertility in China, particularly the roles of the one-child policy, social insurance, and parental expectations, and we used the model to predict impacts of a number of policy changes designed to

reduce the high ratio of boys to girls at birth.

Our model provides a rough match to sex ratios of selected groups in subsamples of the 1982, 1990 and 2000 censuses from China. However we did not attempt to undertake a systematic calibration using a wider set of data as the three period model studied in this paper is rather simplified and stylized along a number of dimensions and this makes it hard to directly confront all of the predictions of the model to the data. Instead, the focus of the analysis has been on using this simple model to gain qualitative insights of the impacts of a number of different policy changes, some that have already occurred, and others that are hypothetical policy options not yet undertaken that might be considered for addressing the sex ratio imbalance problem in China.

We show that moving from a stringent one-child policy to a one-and-half-child policy (second allowed if the first is a girl) would dramatically decrease the sex ratio at birth from 125 to 106. When social insurance coverage is universal, changes in the sex ratio are quite limited if the social attitude remains the same (i.e. parents do not modify their expectation on children's transfer), which suggests that at this moment the underdeveloped social insurance does not contribute to substituting sons' role. However, we also notice that social insurance could play a significant role if social attitude regarding child transfer shifts. Meanwhile, the model predicts that if parents are adaptive and consider the "can-not-marry" risk, the sex ratio will drop from 125 to 110 under the one-child policy, and by a smaller amount if the one-child policy is relaxed.

Our analysis not only models couples' fertility behavior in the context of interactions between birth control policy, social insurance, and parental expectation, but also has policy implications regarding optimal ways to draw the imbalance back to normal.

Although relaxing the one-child policy is the most intuitive option, this could lead to a significant increase in total fertility rate, especially among poor people. The mechanism of social insurance is more complicated, and the change in social attitude could be very slow. Raising sex selection cost like enforcing the abortion ban could prevent people from distorting their fertility choices, but from the eyes of the government, it is very costly to enforce the regulation. However, shifting parental expectations by publicity campaigns and education to let parents be aware of the “marriage squeeze” sons have to face, could dramatically help alleviate the gender imbalance issue. Such a policy is most effective because couples themselves would adjust their fertility choices without serious government interventions. It also avoids the concerns of rising population growth.

There are many extensions to this work that could provide further insights into the gender imbalance problem. First, this paper focuses on the “old-age support” motive for childbearing; the altruism motivation is also important. Modeling both motivations simultaneously could help clarify their relative importance in formulating the son preference. Second, while the three-period framework in this paper enables us to compare the calibrated outcomes with actual ratios for several census years, it lacks the adequate power to predict peoples’ fertility behavior from year to year. A life-cycle model where a period corresponds to one year would be useful to match the actual time series, and to understand the timing of childbearing. Third, this paper calibrates the structural model for various counterfactual experiments on the basis of arbitrary choices over some key parameters, such as sex selection cost and children’s future transfer rate. Although we undertake sensitivity analysis on sex selection cost and endogenize children’s transfer distribution, it would be better if we could utilize micro-level data and empirically estimate these param-

eters by matching moments from simulations to the data. We leave these extensions for future work.

1.7 Appendix

1.7.1 Numerical Root-Finding Procedure

In principle, under the one-child policy, equation $\text{LHS}(s_t) = \text{RHS}(s_t)$ defines an optimal s_t , which in turn yields the optimal value of consumption, transfer and the maximal utility for the middle-aged parents. We now show that under certain parametric restrictions, the solution to $\text{LHS}(s_t) = \text{RHS}(s_t)$ exists and is unique.

By inspection, we see the following four results. First, given the child mortality risk, Equation (1.3) says parents with precautionary saving motives (assuming they can borrow against their pension benefit³⁶) have

$$s_t \geq s^{\text{Min}} = -\frac{\text{SI}_t}{R_{t+1}(1 - \alpha_t)W_t}$$

Second, due to the non-negativity of current parents' consumption (C_t^m) and transfer (d_t),

Equation (1.2) says the maximal saving parents could have³⁷ is

$$s^{\text{Max}} = (1 - b) - \frac{a + c}{(1 - \alpha_t)W_t} - \frac{\phi}{1 - \alpha_t}$$

³⁶If we assume financial market is incomplete, then s^{Min} is zero.

³⁷If grandparents are willing to sacrifice so that $d_t < 0$, we still have the non-negativity of current consumption (C_t^m and C_t^o), and the upper bound of saving will be increased by $(R_t s_{t-1}(1 - \alpha_{t-1})W_{t-1} + \text{SI}_{t-1})/(f_{t-1}(1 - \alpha_t)W_t)$.

Third, for those $s_t < s^{\text{Max}}$, $\frac{\partial \text{LHS}(s_t)}{\partial s_t} > 0$, $\lim_{s_t \rightarrow s^{\text{Max}}} \frac{\partial \text{LHS}(s_t)}{\partial s_t} = +\infty$. Fourth, for those $s_t > s^{\text{Min}}$, we have $\frac{\partial \text{RHS}(s_t)}{\partial s_t} < 0$, and $\lim_{s_t \rightarrow s^{\text{Min}}} \frac{\partial \text{RHS}(s_t)}{\partial s_t} = +\infty$. Therefore, if the parameter values are chosen such that $s^{\text{Min}} < s^{\text{Max}}$, there exists a unique solution to the equation $\text{LHS}(s_t) = \text{RHS}(s_t)$, as displayed in Figure 1.13. However, if parents' income is so low that $(a+c)/((1-\alpha_t)W_t)$ is very large, it is possible that they can afford neither the fixed sex selection price nor the fixed childbearing cost³⁸. In this case, we have $s^{\text{Min}} > s^{\text{Max}}$, so it is not surprising that there is no solution to the equation. Those poor parents will go bankrupt if they decide to have kids.

However, even if there exists a unique solution to the equation $\text{LHS}(s_t) = \text{RHS}(s_t)$, it has no analytical form. Because $\text{RHS}(s_t)$ involves integrals and numerical integration is time-consuming, the process to find s_t^* is excruciatingly slow.

Here, the first time-saving step is to construct discrete approximations³⁹ to the Beta and lognormal distributions as opposed to full-fledged numerical integrations. In this way, we can transform integration operations in Equation (1.11) into summations over values evaluated at each grid point. However, approximating the Beta and the lognormal distribution separately will end up with nested summations and make the whole process time-consuming. For example, under the one-child policy scenario, we will face two double integrals on the RHS and we need to calculate four nested summations.

Therefore, to further speed up the computation, we propose a simplification method

³⁸This is different from the ‘‘Double-Income-No-Kid’’ problem, where the issue is on the income-varying part of sex selection and childbearing.

³⁹An n -point equiprobable grid is created based on Carroll (2011). In terms of approximating the lognormal distribution, we define a set of points on the $[0, 1]$ interval as $\sharp = \{0, 1/n, 2/n, \dots, 1\}$. Denote the inverse of the lognormal distribution as \mathcal{F}_\sharp^{-1} , and define the points $\sharp_i^{-1} = \mathcal{F}_\sharp^{-1}(\sharp_i)$. Then define $\theta_i = \int_{\sharp_{i-1}^{-1}}^{\sharp_i^{-1}} \theta d\mathcal{F}(\theta)$. The θ_i represent the mean values of θ in each of the regions bounded by the \sharp_i^{-1} endpoints, and are used to approximate the true lognormal distribution.

by constructing a discrete approximation to the product of Beta and lognormal distributions so that double integral could be reduced to single integral. In detail, we first create the discrete approximation separately for Beta and lognormal distributions based on 100-point equiprobable grid. Although the parametric form of the combined distribution $d_{t+1}w_{t+1}$ is unknown, its empirical distribution could be simulated by multiplying the grid points from the previous approximations ($100 \times 100 = 10,000$ grid points). Then, we can construct a discrete approximation to this empirical distribution based on an n -point equiprobable grid, in which we treat $d_{t+1}w_{t+1}$ as one single random variable⁴⁰.

To test the speed and accuracy of this simplification method, we pick the “one-girl, without social insurance, myopic parents” scenario as an experiment. For a given combination of $W_{t-1,j}$ and $W_{t+1,k}$, we solve the optimal s_t^* through root-finding for parents earning $[0.1, 6]\overline{W}_t$. For 60 types of parents, it takes around 2.35 minutes to solve s_t^* using separate approximations to Beta and lognormal distributions, while the time is shortened to 0.05 minutes using the above simplification method and choosing $n = 25$, which is almost 50 times faster. In terms of accuracy, the discrepancies in s_t^* between the two methods are within the range $[6.8 \times 10^{-7}, 0.00018]$. This experiment demonstrates that our simplification method helps speed up the whole process without sacrificing accuracies.

However, even after adopting the above simplification method, we still face a double integral in the two-child cases. Inspired by the method of treating the product of two random variables as a single one, we further treat the sum of two random variables as a single

⁴⁰Although d_{t+1} and w_{t+1} are two different random variables, they always enter the model in the format of $d_{t+1}w_{t+1}$, which inspires us to think the feasibility of treating them as one random variable.

one. First, we construct a discrete approximation to the distribution of $d_{t+1}w_{t+1}$. Second, if both kids grow up, the RHS will involve terms like $\sum_{i=1}^2 d_{t+1,i}w_{t+1,i}$. From previous example, our simplification method could find close approximation to $d_{t+1,i}w_{t+1,i}$, which means we can adopt the same method again to approximate $\sum_{i=1}^2 d_{t+1,i}w_{t+1,i}$. In this way, the four-dimensional integral could be simplified into a single integral. We compare the speed and accuracy in solving optimal s_t^* under the “two-girl” scenario using two approximation methods. The time to solve s_t^* for 60 types of parents is only 0.23 minutes, while it takes around 45.7 minutes using separate approximations to Beta and lognormal distributions. What’s more, the discrepancies in s_t^* between the two methods are less than 8.2×10^{-5} . Not surprising, the reduction from four nested sums to one loop of sum tremendously improves the efficiency, for a given tolerance level.

Finally, picking up the starting point is also a subtle issue. On the one hand, since both $\text{LHS}(s_t)$ and $\text{RHS}(s_t)$ have continuous derivatives, Newton’s methods are sufficient. On the other hand, savings are determined by both precautionary motive and intertemporal consumption smoothing motive. Thus, the starting point of the numerical root-finding procedure should be close to s^{Min} , which is determined by $\text{RHS}(s_t)$, rather than s^{Max} , which is derived by minimizing current consumption. This works well in all our cases.

1.7.2 The Solution for Two-Child Models

The model for the “two boys with sex selection, social insurance and adaptive parents” scenario is described as Equation (1.13)-(1.18). The first-order conditions will be the same as Equation (1.7) and (1.8). We can substitute out d_t , which leaves us an equation on s_t as follows:

$$\begin{aligned}
\text{LHS}(s_t) &= \frac{1 + \eta}{(1 - \alpha_t)(1 - s_t - 2b)W_t - (2a + c) - \phi W_t + (R_t s_{t-1}(1 - \alpha_{t-1})W_{t-1} + \mathbf{SI}_{t-1}) / f_{t-1}} \\
\text{RHS}(s_t) &= \delta R_{t+1} \mathbb{E}_t \left(\frac{1}{(d_{t+1,1} + d_{t+1,2})w_{t+1}(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + \mathbf{SI}_t} \right) \\
&= \frac{\delta p_1 p_2 R_{t+1}}{R_{t+1}s_t(1 - \alpha_t)W_t + \mathbf{SI}_t} \\
&+ \int_{\mathbb{R}} \frac{\delta p_1(1 - p_2)\mathfrak{K}_2 R_{t+1}}{e^{y_2}(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + \mathbf{SI}_t} f_{Y_2}(y_2) dy_2 \\
&+ \int_{\mathbb{R}} \frac{\delta p_1(1 - p_2)(1 - \mathfrak{K}_2) R_{t+1}}{e^{x_2}(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + \mathbf{SI}_t} f_{X_2}(x_2) dx_2 \\
&+ \int_{\mathbb{R}} \frac{\delta(1 - p_1)\mathfrak{K}_1 p_2 R_{t+1}}{e^{y_1}(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + \mathbf{SI}_t} f_{Y_1}(y_1) dy_1 \\
&+ \iint_{\mathbb{D}} \frac{\delta(1 - p_1)\mathfrak{K}_1(1 - p_2)\mathfrak{K}_2 R_{t+1}}{(e^{y_1} + e^{y_2})(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + \mathbf{SI}_t} f_{Y_1}(y_1) f_{Y_1}(y_2) dy_1 dy_2 \\
&+ \iint_{\mathbb{D}} \frac{\delta(1 - p_1)\mathfrak{K}_1(1 - p_2)(1 - \mathfrak{K}_2) R_{t+1}}{(e^{y_1} + e^{x_2})(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + \mathbf{SI}_t} f_{Y_1}(y_1) f_{X_2}(x_2) dy_1 dx_2 \\
&+ \int_{\mathbb{R}} \frac{\delta(1 - p_1)(1 - \mathfrak{K}_1) p_2 R_{t+1}}{e^{x_1}(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + \mathbf{SI}_t} f_{X_1}(x_1) dx_1 \\
&+ \iint_{\mathbb{D}} \frac{\delta(1 - p_1)(1 - \mathfrak{K}_1)(1 - p_2)\mathfrak{K}_2 R_{t+1}}{(e^{x_1} + e^{y_2})(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + \mathbf{SI}_t} f_{X_1}(x_1) f_{Y_2}(y_2) dx_1 dy_2 \\
&+ \iint_{\mathbb{D}} \frac{\delta(1 - p_1)(1 - \mathfrak{K}_1)(1 - p_2)(1 - \mathfrak{K}_2) R_{t+1}}{(e^{x_1} + e^{x_2})(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + \mathbf{SI}_t} f_{X_1}(x_1) f_{X_2}(x_2) dx_1 dx_2
\end{aligned}$$

where we have

$$\forall z \in \{y_1, x_1, y_2, x_2\}, z \equiv \text{Log}(d_{t+1}^z w_{t+1})$$

$$d_{t+1}^z \sim \text{Beta}(\alpha^z, \beta^z)$$

$$\log(w_{t+1}) \sim \mathcal{N}(\mu, \sigma^2)$$

$$\mathbb{R} \equiv (-\infty, +\infty)$$

$$\mathbb{D} \equiv \mathbb{R} \times \mathbb{R}$$

1.7.3 Benchmark Parameters

Table 1.6 provides a summary of parameters for the benchmark calibration. First, since we solve a three-period life-cycle model in a partial equilibrium framework, we need to set values for interest rates, wage levels, etc. across periods. (1) R_t is computed as the product of annual real gross interest rate between 1986 and 2005, which is slightly smaller than 1. One might conjecture that the real interest rate will stay at a relatively low level between 2006-2025, so we assume $R_{t+1} = (1 + 1\%)^{20}$. (2) g_t and g_{t+1} represent the real GDP growth rate, which are used to approximate the real wage growth. The geometric average annual GDP growth rate for period 1986-2005 is 9.56%. We pick 6% as its annual growth during 2006-2025, accounting for the recent financial crisis. (3) $\bar{W}_{t+s}, \forall s \in \{-1, 0, 1\}$ are the average labor income. We normalize $\bar{W}_t = 1$ and set $\bar{W}_{t-1} = \bar{W}_t/g_t$ and $\bar{W}_{t+1} = g_{t+1}\bar{W}_t$. (4) Within each period, we define relative wage ($w_{t+s,i}, \forall s \in \{-1, 0, 1\}$) as the ratio of individual wage to average wage level. There is a cross-sectional distribution of this relative wage. $w_{t,i}$ and $w_{t-1,i}$ follow the empirical income distribution for year 2005 and 1985 correspondingly. $w_{t+1,i}$ is assumed to follow

Table 1.6: Benchmark Parameters

Parameter	Symbol	Value
<i>Macro Economy</i>		
Number of years per period	T	20
Risk-free real interest rate 1986-2005	R_t	0.940
Risk-free real interest rate 2006-2025	R_{t+1}	1.220
GDP growth rate 1986-2005	g_t	5.210
GDP growth rate 2006-2025	g_{t+1}	2.207
Average labor income at period t	\overline{W}_t	1.000
<i>Preferences</i>		
Discount factor	δ	0.818
Parents altruism towards grandparents	η	0.250
<i>Child Investment</i>		
Mortality rate between age 0 and 19	p	0.025
Childbearing cost: fixed part	a	0.050
Childbearing cost: income-varying part	b	0.100
Sex selection cost: fixed part	c	0.050
Sex selection cost: income-varying part	ϕ	0.050
Mean transfer from married sons (as a share of income)	m_b	18%
Mean transfer from daughters (as a share of income)	m_g	10.5%
Discount factor for mean transfer from unmarried sons	λ	0.75
<i>Social Insurance</i>		
Marginal tax rate	α_t	0.080
Actuarial fairness index	β_t	0.800
Minimal pension ratio	γ_t	0.100
<i>State Variables</i>		
Total fertility rate in 1985	f_{t-1}	2.600
Average private saving rate in 1985	s_{t-1}	0.150

the same distribution as $w_{t,i}$ under the *Raw* weighting matrix. The *Adj* weighting matrix assumes an income-varying expectation on such distribution as discussed in section 1.4.2.

Second, the utility function is set as the natural logarithm, and the intertemporal discount factor is set as $\delta = 0.99^{20}$. In addition we also need a value for η (parents' attention on grandparents' utility). As expected, there is no direct data measuring this, and it should vary across households. Boldrin et al. (2005) calibrates their model using

Table 1.7: Relative Income Distribution for Calibration

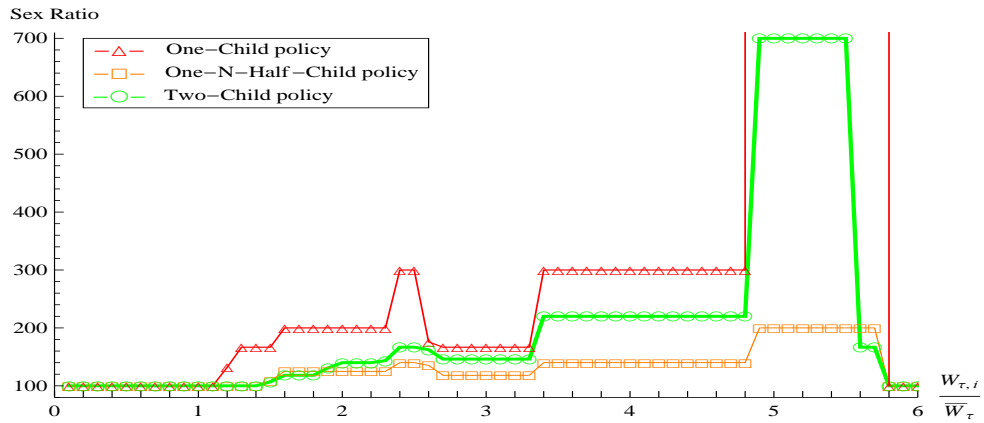
Percentile	0-5%	6-10%	11-20%	21-40%	41-60%	61-80%	81-90%	91-100%
w_{t-1}	0.4629	0.5528	0.6480	0.7844	0.9500	1.1563	1.4077	1.9196
w_t	0.2414	0.3553	0.4595	0.6340	0.8733	1.2010	1.6507	2.7593
Means of w_{t+1}	0.2414	0.3553	0.4595	0.6340	0.8733	1.2010	1.6507	2.7593

$\eta = 0.185$ for England. Here we set it as 0.25, since influenced by the Confucian doctrine, Chinese people tend to care more about their elderly parents' old-age life.

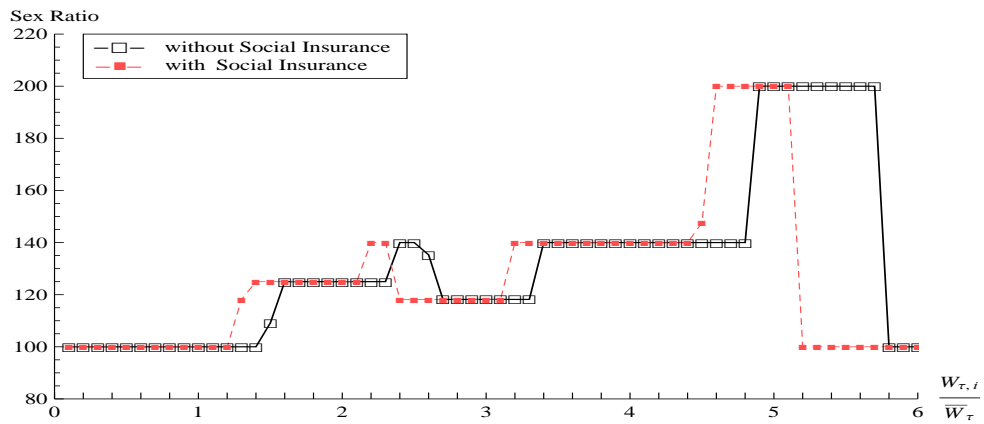
Third, child mortality rate between age 0 and 19 is calculated based on the life table from the World Health Organization.

Fourth, social insurance benefit reflects a combination of minimum living standards and personal contribution history. The minimum pension is a small percentage of average wage level, and the contribution-based benefit is positively related to individual contribution. From year 2000 onwards, Chinese workers on average contribute 8% of their wages to the social insurance trust fund. We further assume the actuarial fairness index equal 0.8, and the minimum pension ratio is 10%. We also assume that grandparents in this model are not covered by any pension system.

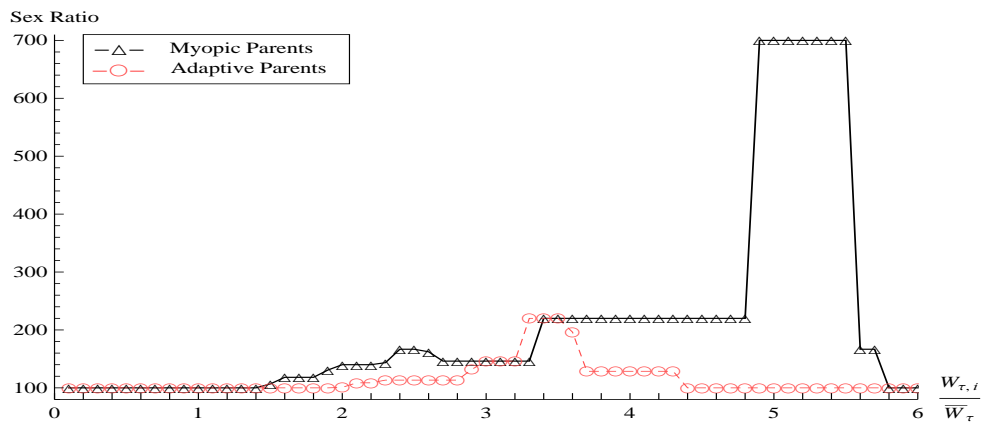
Finally, values for the two state variables are chosen to be TFR and the saving rate of the middle-aged in the 1980s, where we approximate the latter by household saving rate in Modigliani and Cao (2004).



(a) Without Social Insurance, Myopic Parents

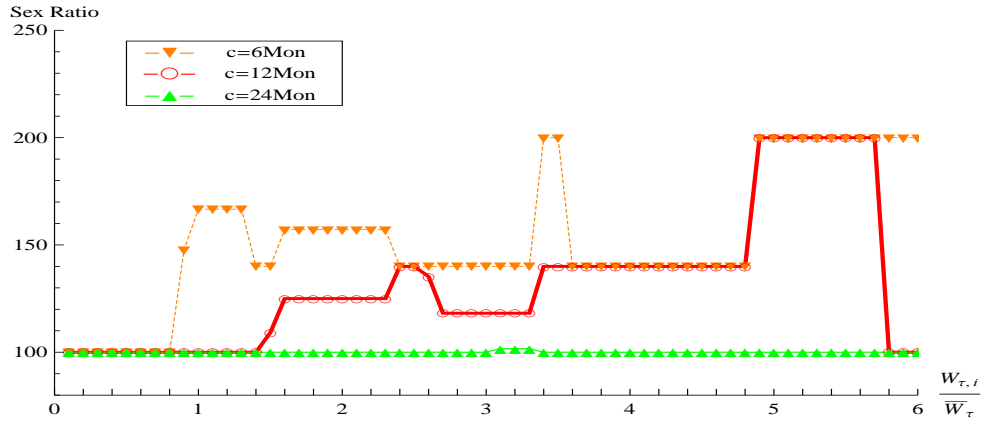


(b) One-N-Half-Child Policy, Myopic Parents




(c) Two-Child Policy, With Social Insurance

Figure 1.10: Sex Ratio by Income




(a) One-N-Half-Child Policy, Without Social Insurance, Myopic Parents




(b) One-N-Half-Child Policy, With Social Insurance, Adaptive parents

Figure 1.11: Sex Ratio under Different Fixed Cost c



(a) One-N-Half-Child Policy, Without Social Insurance, Myopic Parents



(b) One-N-Half-Child Policy, With Social Insurance, Adaptive Parents

Figure 1.12: Sex Ratio under Different Income-Varying Cost ϕ




Figure 1.13: Numerical Root-Finding: An Example

Table 1.8: Society-Wide Sex Ratio under Different Choices of Fixed Cost c

Myopic Parents		6 Months		12 Months		24 Months	
		<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>
SI-No	One-Child	249.1	337.8	113.7	135.6	100.0	100.0
	One-N-Half-Child	120.3	131.7	102.5	108.8	100.0	100.0
	Two-Child	137.0	172.7	103.1	112.9	100.0	100.0
SI-Yes	One-Child	393.8	433.0	119.4	147.7	100.0	100.0
	One-N-Half-Child	142.8	150.4	102.6	109.0	100.0	100.0
	Two-Child	175.5	210.9	104.4	115.5	100.0	100.0
Adaptive Parents		6 Months		12 Months		24 Months	
		<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>
SI-No	One-Child	135.1	185.8	104.0	115.3	100.0	100.0
	One-N-Half-Child	109.0	119.7	100.8	103.9	100.0	100.0
	Two-Child	110.7	130.2	100.6	102.9	100.0	100.0
SI-Yes	One-Child	175.3	214.5	105.2	117.8	100.0	100.0
	One-N-Half-Child	116.6	125.8	100.8	104.0	100.0	100.0
	Two-Child	121.3	141.8	101.9	109.2	100.0	100.0

Table 1.9: Society-Wide Sex Ratio under Different Choices of Income-Varying Cost ϕ

Myopic Parents		9 Months		12 Months		15 Months	
		<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>
SI-No	One-Child	141.8	195.9	113.7	135.6	105.1	117.4
	One-N-Half-Child	112.2	124.4	102.5	108.8	100.8	103.9
	Two-Child	117.5	141.3	103.1	112.9	100.6	102.9
SI-Yes	One-Child	177.2	235.1	119.4	147.7	105.2	117.8
	One-N-Half-Child	112.6	124.8	102.6	109.0	100.8	104.0
	Two-Child	124.1	151.7	104.4	115.5	100.6	103.0
Adaptive Parents		9 Months		12 Months		15 Months	
		<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>
SI-No	One-Child	110.4	136.6	104.0	115.3	100.0	100.0
	One-N-Half-Child	103.4	110.9	100.8	103.9	100.0	100.0
	Two-Child	105.1	117.9	100.6	102.9	100.0	100.0
SI-Yes	One-Child	122.5	157.0	105.2	117.8	100.0	100.0
	One-N-Half-Child	104.0	111.7	100.8	104.0	100.0	100.0
	Two-Child	107.1	124.0	101.9	109.2	100.0	100.0

Table 1.10: Simulated Equilibrium Distribution of Children's Transfer Rate

		Myopic		Adaptive	
		$\eta = 0.25$	$\eta = 0.5$	$\eta = 0.25$	$\eta = 0.5$
		Mean (S.E.)	Mean (S.E.)	Mean (S.E.)	Mean (S.E.)
SI-No	One-Child	11.47 (1.38)	22.10 (2.00)	11.46 (1.39)	22.10 (2.02)
	One-N-Half-Child	11.04 (1.68)	20.99 (2.18)	11.03 (1.68)	21.03 (2.12)
	Two-Child	10.63 (2.18)	19.94 (3.06)	10.64 (2.19)	19.96 (3.01)
SI-Yes	One-Child	15.08 (2.02)	28.92 (3.50)	15.08 (2.02)	28.90 (3.51)
	One-N-Half-Child	14.75 (1.47)	27.46 (2.08)	14.72 (1.47)	27.43 (2.10)
	Two-Child	14.38 (1.26)	26.36 (1.48)	14.40 (1.28)	26.35 (1.48)

Table 1.11: Sex Ratio Based on Simulated Equilibrium Distribution of Children's Transfer Rate

		Myopic				Adaptive			
		$\eta = 0.25$		$\eta = 0.5$		$\eta = 0.25$		$\eta = 0.5$	
		<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>
SI-No	One-Child	112.5	127.0	104.5	116.8	103.3	107.0	100.0	100.0
	One-N-Half-Child	102.0	106.6	101.5	105.4	100.0	100.0	100.0	100.0
	Two-Child	103.0	112.6	101.0	103.8	100.0	100.0	100.0	100.0
SI-Yes	One-Child	119.4	147.7	100.0	100.0	105.2	117.8	100.0	100.0
	One-N-Half-Child	102.6	109.0	100.0	100.0	100.8	104.0	100.0	100.0
	Two-Child	100.6	103.0	102.1	104.5	100.0	100.0	100.0	100.0

Table 1.12: Calibrated Sex Ratio for Different Channels of Social Insurance

		Myopic		Adaptive	
		<i>Raw</i>	<i>Adj</i>	<i>Raw</i>	<i>Adj</i>
SI-No	One-Child	113.7	135.6	104.0	115.3
	One-N-Half-child	102.5	108.8	100.8	103.9
	Two-Child	103.1	112.9	100.6	102.9
SI-Yes Benchmark	One-Child	119.4	147.7	105.2	117.8
	One-N-Half-child	102.6	109.0	100.8	104.0
	Two-Child	104.4	115.5	101.9	109.2
SI-Yes Benchmark + Liquidity Constraint	One-Child	119.4	147.7	105.2	117.8
	One-N-Half-Child	100.9	104.2	100.0	100.0
	Two-Child	102.1	110.0	100.0	100.0
SI-Yes Benchmark + Social Attitude Change	One-Child	105.7	119.0	100.0	100.0
	One-N-Half-Child	100.9	104.2	100.0	100.0
	Two-Child	100.7	103.1	100.0	100.0
SI-Yes Benchmark + Liq. Cons. + Δ Soc. Attitude	One-Child	105.7	119.4	100.0	100.0
	One-N-Half-Child	100.0	100.0	100.0	100.0
	Two-Child	100.0	100.0	100.0	100.0

Chapter 2

Risky Child Investment, Fertility, and Social Insurance in China

2.1 Introduction

2.1.1 Motivation

As the most populous country, China has had a continuous decline in its total fertility rate (TFR) since the 1960s. A nation-wide family control policy that restricts urban families to have only one child and rural families to have at most two kids was implemented in the 1970s to control the population growth, and it appears to be the driving force behind the TFR decline. However, the time-series data show that TFR in China peaked at 5.9 kids in 1967, dropped dramatically to 3.1 in 1977 *before* the start of the one-child policy, then fell between 1978 and 1980 from 2.9 to 2.6. The TFR remained around 2.6 until 1986, and then entered another period of decline from 2.6 to 1.8 between 1987 to 1998. After that, the TFR stabilized around 1.8 until 2009.

When the one-child policy was initially implemented, it prevented people who already have one or two kids from bearing another one and restricted people without children to have only one child. The sharp decline in TFR between 1978 and 1986 comes from the reduction of childbearing from those who already have children. However, the decline after 1986 is different: more and more people decided not to have children instead of cutting the number of children they wanted to have because of the policy constraint.

Thus, it is clear that the one-child policy can not explain the big drop from 1967 to 1977 prior to its imposition, and it is not strong enough to explain the further drop between 1986 and 2009 when it is already widely publicized and enforced. In fact, especially since 2000, the one-child constraint has been relaxed in some places in China. Therefore, it seems like a puzzle why the TFR kept declining after 1985. The purpose of this paper is to investigate whether the advent of social insurance can explain the decline in TFR between 1986 and 2009.

Existing literature demonstrates a significantly negative correlation between TFR and social insurance generosity. China launched its social insurance program in the early 1990s, which motivates us to examine whether the introduction of this program was the primary cause of the decline in TFR. Although the development of social insurance coincides with the drop in TFR, this observation does not reveal by what means this program was able to affect peoples' fertility choices, since many other factors also changed during the same period. We focus on several channels through which social insurance affects childbearing decisions and show that this program helps reduce TFR particularly through the liquidity constraints social insurance tax rates induce on the poor, while increases TFR if people can borrow against their pension benefits through informal channels. Thus, factors outside the one-child policy and social insurance may play a dominating role in reducing TFR and keeping it at a very low level after 1986; we acknowledge that other effects that affect return on investment in children could be contributing factors.

2.1.2 International Background

According to US Census Bureau, world population is expected to reach more than 9 billion in 2050, accompanied with a falling growth rate, which was more than 2% in the 1960s and is projected to be only 0.5% in 2050. However, the pattern is quite uneven across the globe: the population growth will be contributed exclusively by the developing countries; while the population in the developed nations will most likely decline because of the total fertility rate falling below the replacement levels. For example, nowadays countries in Middle East and Sub-Sahara Africa still experience a growth as high as 3%; while the growth rate is less than 1% in developed countries and even becomes negative in Central and Eastern Europe.

Given that one major motive for childbearing is the old-age support, one might ask whether the development in social insurance program, which provides an alternative for retirement funding, can explain the sharp differences in population growth between developed and developing countries. Indeed, under any current classification standard¹, there is much difference in social insurance coverage between developed and developing nations. Within the former group, social insurance was established several decades ago, has developed into a rather comprehensive system, and is run extremely proficiently. In the US, several government-sponsored programs are well designed to take care of the post-retirement lives, like Social Security, Medicare, the Pension Benefit Guaranty Corporation (PBGC) program, the railroad retirement program, etc. And in the case of Social

¹To define developed and developing countries, IMF and World Bank usually look at two criteria: (1) income per capita, (2) industrialization. United Nations Development Programme devised a human development index (HDI) that considers three dimensions: (1) life expectancy, (2) years of schooling, and (3) GNI per capita (PPP US\$).

Security, all administrative duties are performed at a cost of 0.9 percent of total expenditures, or less than a penny per dollar². Similar retirement funding programs are offered in all OECD countries, and pension benefits are very generous in welfare states like Sweden. Within the latter group, many developing countries are in the process of setting up a formal social insurance program or undergoing dramatic reforms. For example, Chile introduced a system of privately managed individual accounts in 1981, and this Chilean model was adopted by 10 other countries in Latin America since 1990. Despite a series of improvements made to the system, however, it still faced several challenges, including high administrative fees. In 2008, the system was overhauled. According to the international research program at the U.S. Social Security Administration, most developing countries only place the old-age security program in their priority list, but not the unemployment or health insurance; in some Asian nations, only certain type of people can apply for unemployment or health insurance benefits; and in most African countries, programs of similar type are of limited scope.

In the context of comparison between developing and developed countries, it seems that there exists a strong correlation between peoples' fertility choices and the social insurance progress. When the social insurance program is comprehensive and generous, people tend to raise fewer kids, while they may rear more with a less-developed program.

Several papers provide empirical tests of the old-age security motivation for child-bearing by looking at the influence of social insurance on fertility in developed countries. Boldrin et al. (2005) argue that an increase in government provided old-age pensions is strongly correlated with a reduction in fertility. They show that the impact of social secu-

²This data is from <http://www.ssa.gov/oact/STATS/admin.html>.

rity on fertility is sizeable in their model, which accounts for 55% to 65% of the observed Europe-US fertility differences both across countries and over time, while the altruism motivation of childbearing is shown to have a small effect.

Billari and Galasso (2008) propose an empirical test using the Italian pension reforms in the 1990s as a natural experiment, which suddenly and substantially decreased pension prospects for a large group of individuals and thus introduced a clear discontinuity in the size of future pension benefits across workers. The empirical results identify a clear and robust effect: less generous future pensions have induced more post-reform fertility, which is again consistent with the old-age security motive.

In addition, Ehrlich and Kim (2007) show that social security taxes and benefits generate incentives to reduce both family formation and fertility. Through a dynamic over-lapping generation model and using panel data from 57 countries over 32 years, they illustrate that the social security tax measures account for a non-trivial part of the downward trends in fertility; and the result is significant and sizeable for OECD countries, but insignificant for non-OECD countries. Similarly, Zhang and Zhang (2004) find that the estimated coefficient on social security is significantly negative in the fertility equation, using data on a group of market economies in the 1960-2000 period.

To sum up, for the seemingly negative correlation between social insurance progress and peoples' fertility choices, the above literature shows that this is true at least in developed countries. They further proceed to interpret it as a causality relationship.

Finally, it is worthwhile to mention the often neglected feedback channel: while social insurance is *associated* with a lower fertility, the prospect of fewer kids poses a solvency challenge for a pay-as-you-go system. In this sense, the generosity of a pension

program itself is inadvertently contributing to its own funding problem in the long run.

2.1.3 Contribution

We develop a tractable life-cycle model to capture peoples' decisions on whether or not to have children and how many they would like to have. In light of the old-age support motive, in our model parents view children as risky investment goods, where the risks come from mortality that children can not grow up to adults and the expected transfer parents obtain in the next period. The major feature of our framework is a macro aggregation based on micro optimization. Instead of focusing on a representative agent in the general equilibrium framework, we solve the individual dynamic fertility choice model in a partial equilibrium framework for agents with heterogeneous budget constraints, then aggregate individual fertility choices to obtain the society-wide TFR.

Our calibration results match the declining trend in TFR quite well. A decomposition of the calibrated TFRs reveals that the combined price and income effects of social insurance tend to increase the TFR by a limited amount, while imposing the liquidity constraint helps reduce the fertility rate significantly. Moreover, the price and income effects are heterogeneous depending on peoples' positions in the income distribution: relatively poor people tend to have more kids since the two effects work in the same direction, while income effect dominates the price effect for relatively rich people so that their fertility rate falls. Our results are not consistent with the negative relationship between social insurance tax rate and fertility rate in Boldrin et al. (2005). An increase in tax rate has trivial impact on peoples' fertility choices in our framework. Beyond this, we bring out

other factors that could have important influence on fertility choices such as the expected GDP growth rate, childbearing cost, and expected children's transfer rate. All these factors jointly determine the rate of return on child investment, which we think is the key element people consider when making fertility decisions.

The remaining of the paper is organized as follows: Section 2.2 describes the dynamic fertility choice model and discusses several channels through which social insurance can affect fertility. Section 2.3 presents the main calibration results. In order to examine the relationship between social insurance and fertility, we calibrate our theoretical model to match the 1985 and 2005 TFRs in China, which represents the scenarios with and without social insurance. We further simulate our model under hypothetical scenarios to isolate the role of social insurance from other factors. In Section 2.4, we show how the price and income effects as well as the liquidity constraint affect the fertility choices of families with heterogeneous income. Section 2.5 provides sensitivity analysis on childbearing cost, potential change in social attitudes and the social insurance tax rate. We make some concluding remarks in Section 2.6.

2.2 The Model

The theoretical framework in this paper is a dynamic fertility choice model, which examines peoples' decisions on whether or not to have children and how many children they would like to have. Instead of assuming the number of children is a continuous variable, we suppose peoples' fertility choice is discrete in that they can choose to have

one, two or three children³. Considering that the one-child policy may not be the driving force for the consistent TFR decline after 1985, we do not impose the one-child limit in our framework. To determine the optimal number of children to bear, parents will compare the expected life-time utility under each situation, assuming other decisions like consumption, saving and transferring to elderly parents are made optimally.

2.2.1 A Generic Model

We assume individuals can live three periods: young, middle-age, and old⁴. Young individuals simply consume parents' resources to grow up. When they are middle-aged, they supply one unit of labor, obtain income⁵ (W_t) and make optimal decisions on fertility (0, 1, 2, or 3), consumption (C_t^m), saving (s_t), and transfer to their elderly parents (d_t). Under the old-age support motivation, parents view children as investment goods and expect to receive transfers from them when parents become old next period. The level of attention adult children give to their elderly parents is denoted by η , which can also be regarded as the degree of altruism of adult children (i.e. parents in our model) to their elderly parents (i.e. grandparents). In addition, the cost of rearing children consists of

³Although we observe some families having more than three kids, the maximum number of children considered in our framework is three. Based on the United Nations data, worldwide average TFR was around 4.9 in the 1950s, decreased to 3.5 in the 1980s, 2.6 in the 2000s, and is projected to fall to 2.2 in 2050. In the US, TFR fell from 3.7 in 1960 to 2.1 in 2008; in China this number decreased from 5.5 in 1960 to 1.8 in 2008; and for European Union countries, the TFR was only 1.6 in 2008. Considering all these historical data and the predicted trend, we believe a fertility quota of three children per family is reasonable and should not distort our main conclusion.

⁴To avoid confusion, we call them children, parents (or adult children), and grandparents (or elderly parents).

⁵Here is a rule for notation: lower-case letters usually represent the percentage or ratio, while upper-cases represent the absolute level. For example \bar{W}_t is the society-wide average income, $W_{t,i}$ is the income for individual i , and the corresponding relative income is $w_{t,i} = W_{t,i}/\bar{W}_t$. Similarly, $s_{t,i}$ represents the private saving rate for individual i , while $S_{t,i} = s_{t,i}W_{t,i}$ is his/her private saving amount. For ease of notation, we omit the subscript i .

two parts: a fixed cost (a) and an income-varying cost (bW_t). When individuals become old, they retire and their consumption is financed from three sources: the transfers from their adult children (D_{t+1}), private saving from the previous period ($R_{t+1}s_tW_t$) and social insurance benefits (SI_{t+1}) if they are covered by the program.

In this framework, children are viewed as risky investments in the eyes of their parents. On the one hand, with probability p , young individuals can not grown up (i.e. mortality risk); on the other hand, children's transfers in the next period (D_{t+1}) are uncertain. In detail, we decompose the transfer amount as: $D_{t+1} = d_{t+1}w_{t+1}\bar{W}_{t+1}$, where both transfer rate (d_{t+1}) and children's relative income (w_{t+1}) are random variables⁶. We further assume that d_{t+1} follows a Beta distribution which guarantees the transfer rate is within the range of $[0, 1]$; and w_{t+1} is log-normally distributed, which ensures that the relative income is positive, and the median income is smaller than the mean, as consistent with empirical income distributions.

A generic version of the individual optimization problem when social insurance program is in place is as follows⁷:

⁶We do not assume aggregate growth uncertainties, so \bar{W}_{t+1} will grow in a deterministic way.

⁷In our framework, we don't consider the economies of scale in raising children so that the cost of bearing two kids is twice as much as that of having one child. In addition, we assume transfers from children are independent from each other and follow the same product distribution.

$$\text{Max}_{\{s_t, d_t, f_t\}} \mathbb{U}(C_t^m, C_{t+1}^o, C_t^o) = \{u(C_t^m) + \delta \mathbb{E}_t u(C_{t+1}^o) + \eta u(C_t^o)\} \quad (2.1)$$

s.t.

$$C_t^m = (1 - s_t - d_t)(1 - \alpha_t)W_t - (a + b(1 - \alpha_t)W_t)f_t \quad (2.2)$$

$$C_{t+1}^o = \sum_{i=1}^{f_t} d_{t+1,i} w_{t+1,i} (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \mathbf{SI}_t \quad (2.3)$$

$$C_t^o = f_{t-1} d_t (1 - \alpha_t) W_t + R_t s_{t-1} (1 - \alpha_{t-1}) W_{t-1} + \mathbf{SI}_{t-1} \quad (2.4)$$

$$\mathbf{SI}_t = (\gamma_t \bar{W}_t + \beta_t \alpha_t (1 - b f_t) W_t) R_{t+1} \quad (2.5)$$

$$\mathbf{SI}_{t-1} = (\gamma_{t-1} \bar{W}_{t-1} + \beta_{t-1} \alpha_{t-1} (1 - b f_{t-1}) W_{t-1}) R_t \quad (2.6)$$

$$f_t \in \{0, 1, 2, 3\} \quad (2.7)$$

For fixed f_t the F.O.C.s are:

$$u'(C_t^m) = \delta R_{t+1} \mathbb{E}_t \left(u'(C_{t+1}^o) \right) \quad (2.8)$$

$$u'(C_t^m) = \eta f_{t-1} u'(C_t^o) \quad (2.9)$$

Using logarithmic utility function and after some algebraic manipulation of (2.8) and (2.9), we can substitute out d_t as:

$$d_t = \frac{\eta}{1 + \eta} \left(1 - s_t - b f_t - \frac{a f_t}{(1 - \alpha_t) W_t} \right) - \frac{R_t s_{t-1} (1 - \alpha_{t-1}) W_{t-1} + \mathbf{SI}_{t-1}}{(1 + \eta) f_{t-1} (1 - \alpha_t) W_t} \quad (2.10)$$

The solution is characterized by an equation on s_t : $\text{LHS}(s_t) = \text{RHS}(s_t)$, where

$$\text{LHS}(s_t) = \frac{(1 + \eta) f_{t-1}}{\left(1 - s_t - b f_t - \frac{a f_t}{(1 - \alpha_t) W_t} \right) (1 - \alpha_t) W_t f_{t-1} + R_t s_{t-1} (1 - \alpha_{t-1}) W_{t-1} + \mathbf{SI}_{t-1}} \quad (2.11)$$

However, $\text{RHS}(s_t)$ depends on the number of children parents decide to have. We first show that for $f_t = 1$:

$$\begin{aligned}
\text{RHS}(s_t) &= \delta R_{t+1} \mathbb{E}_t \left(\frac{1}{d_{t+1} w_{t+1} (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \text{SI}_t} \right) \\
&= \frac{\delta R_{t+1} p}{R_{t+1} s_t (1 - \alpha_t) W_t + \text{SI}_t} \\
&+ \int_{\mathbb{R}} \int_0^1 \frac{\delta R_{t+1} (1 - p)}{d_{t+1} w_{t+1} (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \text{SI}_t} \\
&\quad f(d_{t+1}) f(w_{t+1}) d(d_{t+1}) d(w_{t+1}) \tag{2.12}
\end{aligned}$$

where,

$$\begin{aligned}
d_{t+1} &\sim \text{Beta}(\alpha_d, \beta_d) \\
\log(w_{t+1}) &\sim \mathcal{N}(\mu, \sigma^2) \\
\mathbb{R} &\equiv (-\infty, +\infty)
\end{aligned}$$

Double integrations are necessary in eq. (2.12), because we assume both transfer rate d_{t+1} and relative wage w_{t+1} are random variables. Given our parametric assumptions, there is no analytical solution for the equation $\text{LHS}(s_t) = \text{RHS}(s_t)$. The justification on the existence and uniqueness of the optimal solution can be found in section 1.7.1, except that due to different model specifications, the expression for s^{Max} is $1 - bf_t - \frac{af_t}{(1-\alpha_t)W_t}$.

In order to find numerical solutions, we could construct discrete approximations to the two continuous distributions, i.e. find the appropriate grid points to represent the corresponding distributions and approximate the integral by summations. However, approximating the Beta and the log-normal distribution separately will end up with nested

summations and be time-consuming⁸. A more efficient alternative is to approximate the product distribution. Although there is no analytical density function of the product distribution, we could generate a representative sample from it, by first simulating two large samples from the Beta and log-normal distributions separately, and second multiplying the simulated grid points from the above two samples. Denoting $z = d_{t+1}w_{t+1}$, we have reduced the double integral into a single integral⁹ so that RHS for the one-child case can be rewritten as:

$$\begin{aligned} \text{RHS}(s_t) &= \frac{\delta R_{t+1}p}{R_{t+1}s_t(1 - \alpha_t)W_t + SI_t} \\ &+ \int_{\mathbb{R}^+} \frac{\delta R_{t+1}(1 - p)}{z(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + SI_t} f_Z(z) dz \quad (2.13) \end{aligned}$$

After this, we could select a certain number of grid points from the simulated product distribution $f_Z(z)$, and then approximate the integral by summation.

Similarly, $\text{RHS}(s_t)$ for the two- and three-child cases are as follows¹⁰:

$$\begin{aligned} \text{RHS}(s_t) &= \frac{\delta R_{t+1}p^2}{R_{t+1}s_t(1 - \alpha_t)W_t + SI_t} \quad (2.14) \\ &+ \binom{1}{2} \int_{\mathbb{R}^+} \frac{\delta R_{t+1}(1 - p)p}{z_1(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + SI_t} f_Z(z_1) dz_1 \\ &+ \binom{2}{2} \iint_{\mathbb{R}^{+2}} \frac{\delta R_{t+1}(1 - p)^2}{(z_1 + z_2)(1 - \alpha_{t+1})\bar{W}_{t+1} + R_{t+1}s_t(1 - \alpha_t)W_t + SI_t} \prod_{i=1}^2 (f_Z(z_i) d(z_i)) \end{aligned}$$

⁸Applying the discretization to the Beta and the log-normal distribution separately will cause more efficiency problems when people choose to have two or three children. In the two-child case, if both children successfully survive to adulthood, one need to calculate two double-integrals, or do four nested summations. By the same token, the three-child case requires six nested summations to approximate the integrals. Not surprisingly, computing burden will then increase exponentially.

⁹More details on the numerical methods, i.e. discrete approximation and further simplification, are in Section 1.7.1.

¹⁰ $\binom{n}{r}$ represents the number of possible combinations of r objects from a set of n objects, which is calculated as: $\binom{n}{r} = \frac{n!}{r!(n-r)!}$.

$$\begin{aligned}
\text{RHS}(s_t) &= \frac{\delta R_{t+1} p^3}{R_{t+1} s_t (1 - \alpha_t) W_t + \text{SI}_t} & (2.15) \\
+ \binom{1}{3} &\int_{\mathbb{R}^+} \frac{\delta R_{t+1} (1-p) p^2}{z_1 (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \text{SI}_t} f_Z(z_1) dz_1 \\
+ \binom{2}{3} &\iint_{\mathbb{R}^2} \frac{\delta R_{t+1} (1-p)^2 p}{(z_1 + z_2) (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \text{SI}_t} \prod_{i=1}^2 (f_Z(z_i) d(z_i)) \\
+ \binom{3}{3} &\iiint_{\mathbb{R}^3} \frac{\delta R_{t+1} (1-p)^3}{(z_1 + z_2 + z_3) (1 - \alpha_{t+1}) \bar{W}_{t+1} + R_{t+1} s_t (1 - \alpha_t) W_t + \text{SI}_t} \prod_{i=1}^3 (f_Z(z_i) d(z_i))
\end{aligned}$$

2.2.2 Does Social Insurance Really Reduce Fertility Rate?

A close look at the model, however, shows that things are more complicated than the intuitive idea that fertility drops in the presence of social insurance. In fact, social insurance has a heterogeneous effect on fertility and its impact is reflected through several channels: the price effect, the income effect, the liquidity constraint effect, and the change in expectation on future transfers.

The first channel through which social insurance directly affects fertility is the income effect. As an income redistribution program, social insurance favors poor people at the expense of the rich. After the introduction of social insurance, poor peoples' present discounted value of life-time income should increase while that of the rich decreases. Since children, saving, transfer, and consumption are all normal goods, poor people would like to raise more kids and the rich tend to rear fewer.

The second channel is the direct price effect. We assume the income-varying part of the childbearing cost is tax-deductible: parents could either send children to daycare and deduct the child-care expense when filing their tax returns, or reduce their labor supply, like working at home, to take care of their kids. After social insurance tax is introduced,

the actual income-varying cost is $(1 - \alpha_t)bW_t$ instead of bW_t . This is equivalent to a per-child subsidy of the size $\alpha_t bW_t$. Since childbearing becomes cheaper, the fertility rate should increase, for both the poor and the rich.

Combining these two effects, however, we find that the effects work in the same direction for poor people so that they will have more kids; the effects are offsetting for the rich so whether they have more or fewer kids depends on which effect dominates. Since a country's TFR is an average fertility choice of different types of people and given that the income distribution is usually more skewed to the left (far more people below the mean than those above), it could happen that the overall fertility rate could increase after a pension system is introduced, instead of declining as Boldrin et al. (2005) have argued. However, it is too early to conclude that social insurance should increase TFR since there are two indirect channels that this program could affect peoples' fertility decisions.

The third channel is whether people can borrow against their future pension benefits. Our previous argument that social insurance increases the life-time income for the poor but decreases that for the rich holds under the condition that financial markets are complete so that people can borrow against their future pension benefits. This will lead to a negative private saving rate, especially for those poor people whose future pension benefits occupy a large proportion of their life-time income. The assumption of a complete market makes some sense in that several informal channels like borrowing from relatives are available in China. On the other hand, we also need to acknowledge that a mature consumer credit market does not exist in China¹¹. Hence any informal borrowing may be

¹¹This is in great contrast to the advanced economies, where the credit market is well developed. In the US, several state pension benefit agencies have the pension loan programs, and various commercial products are available, like the loan provisions under 401(k) plans.

able to lessen the liquidity constraint, but not eliminate it completely. So the degree of liquidity constraint should be a continuous variable between 0 and 1. This effect will be explicitly considered in our calibration exercises.

The fourth channel is a potential change in peoples' expectations. Once social insurance is introduced, a social attitude on who cares for the elderly may shift. The grown-up children in the next period, after realizing that their elderly parents' retirement lives are taken care of by the pension plans, will feel that their role of supporting elderly parents' old-age life has been substituted, and thus exhibit less altruistic behavior (i.e. a smaller degree of altruism), and in turn transfer less. The current parents, in anticipation of this change, will rationally expect a lower return on child investment, since the distribution of the transfer rate from children moves leftwards, and respond by having fewer kids. This channel will take effect gradually and it is hard to verify whether parents are forward-looking and whether such expectation shifts have happened. We will deal with this issue in the sensitivity analysis part.

In sum, considering all the four possible channels, we see that social insurance may have an ambiguous effect, with the price and income effects pushing up the fertility rate while liquidity constraint and social attitude changes pulling down the TFR.

However we do see in reality that there is an obvious downward trend in fertility over the past few decades. If social insurance's effect is ambiguous, what are other factors that contribute to the decline in TFR? We propose that factors closely related to the rate of return on child investment (i.e. the cost and rewards of childbearing) are responsible. One factor is that as economy develops and living standard rises, the cost of raising children increases. The other factor is that people may predict that the aggregate economic growth

will slow down¹². Both factors imply that children’s value in terms of investment return will decline¹³. Thus, people tend to have fewer kids and allocate extra resources to saving, transfer or consumption.

2.3 Calibration

In this part we calibrate our model under two scenarios (with and without social insurance), to match the TFR in China in 1985 (when social insurance system was absent) and in 2005 (when the government has undertaken a series of structural reforms to the pension system). This 20-year gap¹⁴ also facilitates our overlapping generation framework.

China’s national social insurance program was launched around 1993. Before that, the period 1950-1978 was characterized by two anomalous events—the Great Famine and the 10-year Cultural Revolution. Both have had major impacts on the economy as well as on fertility, especially the Great Famine. Beginning in 1979, the government adopted the “reform and opening-up” policy and the economy experienced dramatic changes in the early 1980s. For all these reasons, we choose the TFR in year 1985 as the empirical target for the scenario without insurance. After the 1993’s launching, various policy reforms took place to the social insurance programs, and by 2005 the system has stabilized.

¹²To have a snapshot on this hypothesis, we find that US real GDP grew at 6.34% in 1935-1939, 5.98% in 1940-1949, 4% during 1950-60s, dropped to 3% in the last three decades of the 20th century, and was only 2.33% between 2000 and 2008. During the same period, the TFR falls from around 3.0 in the 1930s to the current level of 2.06, a decline of 30%.

¹³Based on our assumption that $D_{t+1} = d_{t+1}w_{t+1}\bar{W}_{t+1}$ and the growth in average wage is equal to GDP growth, even if the distribution of d_{t+1} does not shift leftwards, an expected decrease in GDP growth will lead to a smaller \bar{W}_{t+1} , which then lowers D_{t+1} .

¹⁴Throughout the calibration, we assume that a period is 20 years which may be a bit away from the realistic setting, i.e 40 years of working and 20 years of retirement.

Therefore, our calibration target for the scenario with social insurance is the TFR in 2005.

2.3.1 Parameter Choices

For all the parameters, we summarize their values in Table 2.1. Since we calibrate our model to two different years (1985 and 2005), there are four generations in our calibrations, i.e. 1965, 1985, 2005, and 2025.

First, we provide formal justifications for those parameters related to child investment. Childbearing cost consists of two parts: a fixed cost (a) and an income-varying cost (bW_t). For the fixed cost a , we assume it is around 5% of the average income, i.e. one-year's average income. With regard to the income-varying part, Echevarria and Merlo (1999) find that cost to a woman of having a child is about 5% of her working lifetime, while according to Juster and Stafford (1991), hours per week allocated on childcare account for between 6.43% and 18% of parents total available time. In China, children are heavily dependent on parents' support at least before marriage. In order to capture this characteristic, we set b as 10%. In terms of child investment rewards, we assume that children's transfer rate follows a Beta distribution ($d_{t+1} \sim \text{Beta}(\alpha_d, \beta_d)$) and the mean ($\mathbb{E}_t[d_{t+1}]$) is 12.5%.

Second, since we solve our model in a partial equilibrium framework, we need to set appropriate values for interest rates, income levels, etc. across periods. (1) R_{1985} and R_{2005} are computed as the products of annual real gross interest rate for 1966-1985 and 1986-2005¹⁵. And one might conjecture that the real interest rate will stay at a relatively

¹⁵The real interest rates are calculated from nominal rates and inflation rates as $R = 1 + r = (1 + i)/(1 + \pi)$, where r is net real interest rate, π inflation and i nominal rate. The inflation rates are available from 1980 to 2008, and are assumed to be 0 from 1950 to 1979, a time when Chinese economy almost stagnated.

low level between 2006 and 2025, so we assume $R_{2025} = (1 + 1\%)^{20}$. (2) Using historical data, we calculate the real GDP growth rate, and use it to approximate the real income growth between 1966 and 1985 as well as between 1986 and 2005. In addition, we pick 6% as annual GDP growth during 2006-2025, accounting for the recent financial crisis. (3) $\bar{W}_{t+s}, \forall s \in \{-1, 0, 1\}$ are the average labor income: we normalize $\bar{W}_t = 1$ and set $\bar{W}_{t-1} = \bar{W}_t/g_t$ and $\bar{W}_{t+1} = g_{t+1}\bar{W}_t$, where g_t and g_{t+1} are real GDP growth rate for the current and next period. (4) Within each period, we define relative income ($w_{t+s,i}, \forall s \in \{-1, 0, 1\}$) as the ratio of individual income to average level. There is a cross-sectional distribution of this relative income: $w_{1985,i}$ follows the empirical income distribution for year 1985, $w_{2005,i}$ for year 2005. Due to data limitation, $w_{1965,i}$ is assumed to follow the same distribution as $w_{1985,i}$ ¹⁶. Since income distribution is moving slowly and it is difficult for parents to make an accurate conjecture, we assume a static expectation: parents believe that the distribution of the next generation's income will be identical to the income distribution in their own generation. Based on this, $w_{2025,i}$ has the same distribution as $w_{2005,i}$. The details are shown in Table 2.2.

Third, the utility is set as the natural logarithm function; the intertemporal discount factor is set as $\delta = 0.99^{20}$. In addition we also need a value for η (parents' attention on grandparents' utility). As expected, there is no direct data measuring this; and it should vary across countries and even households. Boldrin et al. (2005) calibrate their model using $\eta = 0.185$ for England. Here we set it as 0.5, since under the Confucian doctrine,

The calculated net real interest rate between 1966 and 1985 is around 90%, while it is slightly negative for the period 1986-2005.

¹⁶The earliest income distribution data we could obtain is for year 1995, which is used to approximate the cross sectional distribution of labor income in 1985. Because of a series of disruptions, the income distribution did not change much from 1965 to 1985.

Chinese people tend to care more about their elderly parents' old-age life.

Fourth, we don't have data on child mortality rate between age 0 and 19. However, according to the World Development Indicators database by the World Bank, the mortality under age 5 per 1,000 live births is around 45 in 1995 and 25 in 2005. We conjecture that the mortality rate between age 0 and 19 should be fairly close to the mortality rate under age 5. Therefore, we pick a moderate value 5% for our calibration exercises.

Fifth, social insurance benefit reflects a combination of minimum living expenses and personal contribution history. The minimum pension is a small percentage of average income level and the contribution-based benefit is positively related to individual contribution history. From year 2000 onwards, Chinese workers contributed on average 8% of their incomes to the social insurance trust fund. We further assume the minimum pension ratio is 10%. Due to this setting, a balanced government budget constraint requires that the marginal pension increase should always be less than 1 for one more dollar's tax payment, and thus we set the value of this coefficient to be 0.8. We also suppose middle-age people between 1966 and 1985 were not covered by any pension system, and that the expected social insurance tax rate for 2005-2025 remains at 8%¹⁷.

Finally, in terms of the two state variables—TFR and saving rate in the previous period, we use historical data in the 1960s and 1980s when matching the 1985 and 2005 fertility rates, respectively. In particular, we approximate the saving rate of the middle-age people by household saving rate in Modigliani and Cao (2004).

¹⁷With the rise in the aged dependency ratio, it is expected that the social insurance tax rate might increase in the future. And we will deal with its effect in the sensitivity analysis part.

2.3.2 Calibration Results

Before showing our results, we briefly discuss the calibration procedures. There are three dimensions of heterogeneity here: income of parents $W_{t,i}$, grandparents $W_{t-1,j}$, and the expected income distribution of kids $W_{t+1,k}$. First, we solve individual maximization models with different fertility choices for each combination of (i, j, k) . Second, by comparing the expected life-time utility when $f_t = 0, 1, 2, 3$, we obtain the optimal fertility decision for each income combination pair. Third, we aggregate the individual fertility choices over dimension k , then over j and finally over i , which leads to a society-wide TFR, corresponding to a particular setting on social insurance.

The stylized fact we aim at matching is TFR, which equals 2.60 in 1985's China and 1.80 in 2005, as computed from the World Development Indicators database by the World Bank. Calibrated from the model without social insurance, the weighted average of TFR is 2.55 which is quite close to its empirical counterpart. And the society-wide average private saving rate is 0.56%¹⁸ and people on average transfer 19.76% of their income to elderly parents.

The calibration results for 2005 China is also encouraging. Without applying the liquidity constraint, the calibrated TFR is 1.96 kids with -8.9% private saving¹⁹ and 23.72% transfer rates from adults to their elderly parents when social insurance is in place. When the liquidity constraint is turned on, the calibrated TFR drops to 1.27 kids

¹⁸The calibrated 0.56% saving rate in 1985 is too small as compared to that from Modigliani and Cao (2004) where the saving rate is around 15%. This is because parents in our model invest a lot in children, whose transfers will substitute, to a great extent, the role of saving for retirement. Another reason is our utility setting: the natural logarithm implies a relative risk aversion parameter of 1, while it is commonly estimated to be 3-5 from micro level data. As expected, the more risk averse, the more the households will save.

¹⁹Combining the private saving and the mandatory social insurance saving, the total saving rate is around 11.62%.

with a significant decrease in private saving rate to 6.99% and a relatively stable transfer rate of 22.75%. Although banks normally do not accept pension benefits as collateral, various borrowing channels exist such as from relatives and friends. The liquidity constraint may be binding for some people, but not for others. We find that the average TFR from the models with and without liquidity constraint is around 1.60 kids, which is not far away from the values we observed in the data.

In addition, we find that our calibrations match the declining tendency in TFR and the increasing trend in private saving as Chinese economy develops. However, the decline in TFR seems to contradict our discussion of the price and income effect of social insurance. Hence, we are interested in examining whether our calibration results are consistent with the predictions from our theoretical framework; if so, what are the factors (outside the social insurance program) that contribute to the decline in TFR from 1985 to 2005. We provide a detailed analysis of our calibration results by decomposing and comparing outcomes under various hypothetical environments in the next part.

2.3.3 Analysis of Calibration Results

Comparing the calibrated TFR between 1985 and 2005 can not give us a clear answer on the effect of social insurance on fertility since many other factors also change during this period. To distinguish the impact of social insurance from that of other factors, we calibrate the model under hypothetical environments such as the presence of social insurance in 1985 and the absence of it in 2005. The calibrated results are summarized in Table 2.3. The dark gray cells represent the actual scenarios, while the light gray

ones reflect the hypothetical environments.

First, we suppose parents in 1985 were covered by the social insurance program and calibrate the model using parameter values corresponding to the 1985 economy. It is clear that, other things equal, the combination of the price and income effect of social insurance increases TFR from 2.55 to 2.82, a net change of 0.27, which is consistent with our discussion that increased life-time income and reduced childbearing cost induce people to have more children by running a negative private saving at the current period. However, if people can not borrow against their future pension benefits, calibrated TFR drops from 2.82 to 2.10, indicating a substantial effect of liquidity constraints on reducing fertility.

Second, we assume parents in 2005 were not covered by the social insurance program. The calibrated TFR is 1.80, which can be compared with 1.96 under the scenario with social insurance in 2005. Once again, this shows that the combined price and income effect of social insurance is positive. When we add the liquidity constraint, the TFR falls to 1.27, confirming the substantial effect of this constraint under the 2005 parameter combination.

In sum, Table 2.3 gives us a good example to understand the effect of social insurance on TFR. Although the price and income effects and the impact of liquidity constraint are consistent with our predictions, they, themselves, can not explain the big drop in the calibrated TFR from 1985 to 2005. Looking at Table 2.3 vertically, we see that the 2005 parameter combination generates lower TFRs than the 1985 one, whatever we impose social insurance or not. Factors other than social insurance may play a dominating role in reducing the TFR over time.

As documented in Table 2.1 and 2.2, a set of parameters change their values from 1985 to 2005, including the real interest rate (R_{t+1} and R_t), GDP growth rate (g_t and g_{t+1}), TFR and saving rate in the previous period (f_{t-1} and s_{t-1}), and income distribution of parents (w_t). In order to examine the effect of each factor, we calibrate the model under several hypothetical environments: each time only one parameter changes its value to reflect the 2005 economy, while all the other parameters take their values in 1985. A summary of this comparative statics exercise is reported in Table 2.4.

The first parameter is R_{t+1} , i.e. the real interest rate on private savings between period t and $t + 1$. The change in this parameter value affects the interest income parents will receive when they become old. From Table 2.1, R_{2025} is higher than R_{2005} , indicating a rise in the rate of return on private savings, which will induce a higher private saving rate and a lower fertility. Our result confirms this argument: the TFR drops from 2.55 to 2.49, with an increase in the average saving rate from 0.56% to 1.42%.

The second parameter is g_{t+1} , i.e. the expected GDP growth rate in period $t+1$. This parameter matters because it is used to calculate the society-wide average income \bar{W}_{t+1} , which in turn affects the expected child transfer for everyone ($D_{t+1} = d_{t+1}w_{t+1}\bar{W}_{t+1}$). Historical data shows that the annual real GDP growth rate is more than 9% during the period 1986-2005. However, due to the recent financial crisis, we think it is very hard for China to keep growing at such a high speed for the next 20 years. A less impressive long-run economic growth translates into a lower growth of the average income, and thereby a reduced rate of return on child investment and an expected drop in fertility. Our calibration shows a relatively large drop in TFR from 2.55 to 2.21 kids.

The third parameter we examine is g_t , which is used to compute the average in-

come in the previous period ($\bar{W}_{t-1} = \bar{W}_t/(1 + g_t)$). Since Chinese economy grew at a higher speed during the period 1986-2005 than 1966-1985, the average grandparents' income calculated using g_{2005} is lower than the one using g_{1985} . This means grandparents would become poorer, which would induce a higher transfer. When parents transfer more resources to their elderly parents, fewer resources are left at current period and we expect to see a drop in both saving and fertility. However, since the difference between g_{2005} and g_{1985} is small and it does not affect the rate of return on child investment, the calibrated fertility rate remains roughly the same and the changes in transfer and saving are trivial.

The fourth parameter we are interested in is f_{t-1} , i.e. the state variable indicating the number of siblings of parents. China's TFR in 1965 was 6.08 kids, but this number decreased dramatically to 2.60 in 1985. According to eq. (2.10), where we assume cooperative transfer game among siblings, there is a positive relationship between optimal transfer rate and the number of siblings. Our calibration shows that when f_{t-1} becomes smaller, d_t drops from 19.76% to 19.24%, with TFR unchanged.

The fifth parameter we look at is $R_t s_{t-1}$, which measures grandparents' saving balances. Everything else equal, a higher $R_t s_{t-1}$ represents richer grandparents. Rationally, parents should transfer less to their elderly parents and use the extra resources to save or to raise children. From 1966-1985 to 1986-2005, R_t decreased by a small amount, but the average saving rate increased from 5% to 15%, thus grandparents became richer. Our calibration results reveal that the effect of grandparents' financial position is quite limited in that parents only decrease their transfer by a tiny percentage. Naturally, since $R_t s_{t-1}$ does not directly affect the rate of return on child investment, there is no change in the calibrated TFR.

The final one is income distribution among the current generation of parents, where the gap between poor and rich people enlarges significantly from 1985 to 2005. Our calibration results indicate that TFR decreases while both saving and transfer rates rise. We will have a detailed discussion of the relationship between parents' income position and their fertility decisions in the heterogeneity analysis part below.

2.4 Heterogeneity Analysis

The effect of social insurance is not uniform across the population. In this part we show that the heterogeneous effects of social insurance depend largely on one's position in the income distribution. We focus on the dimension of parents' income (W_t): normalize the average income (i.e. $\bar{W}_t = 1$), and assume relative income ranges from 10% to 6 times of the average. All the other parameter values remain the same as in the above calibrated models.

Different from the benchmark calibration, we only aggregate the individual fertility choices over the dimensions of expected kids' and grandparents' income to obtain a series of weighted average TFRs corresponding to each grid point in parents' income distribution. In addition, we also calibrate our model under a hypothetical scenario, where both parents and grandparents were not covered by the social insurance program in 2005. Thus, a comparison between the outcome from the hypothetical scenario and that from the actual one helps us to single out the impact of social insurance.

2.4.1 Total Fertility Rate

Figure 2.1 describes the optimal fertility choices across parents' income. The blue line is the calibrated TFR for 1985, which can be regarded as a baseline scenario. The green line represents the optimal fertility choices under the hypothetical scenario that both parents and grandparents were not covered by the program in 2005. The orange and the purple lines are results when social insurance is in place. The orange line indicates the price and income effect, while the purple one shows how TFR changes if parents face liquidity constraints.

One common feature across all these scenarios is: starting from the lowest income percentile, fertility rate increases as income rises; however this trend reverses when income is above a threshold level²⁰. The negative relationship between TFR and income when $W_t > \bar{W}_t$, is consistent with previous literature that when parents become rich, the income-varying cost of childbearing becomes higher, which prohibits some of them from raising children. The smaller fertility rate among the poor is because they don't have enough resources to bear the fixed cost of raising children.

Besides the above common feature, there are dramatic differences among these scenarios. First, comparing the blue line with the green one, we find that for every income level, TFR under the hypothetical scenario is lower than that under the baseline scenario, indicating a consistent effect of other factors (such as expected long-run economic growth and interest rate) on reducing TFR from 1985 to 2005. Moreover, the magnitude of the differences between these two scenarios enlarges when parents' income rises, which sug-

²⁰The threshold level is roughly equal to the society average income, although it shifts slightly across different scenarios.

gests that rich people may be more sensitive to changes in other contributing factors. For example, the difference is only 0.05 for people earning $0.2\bar{W}_t$ and less than 1.0 if parents' income is within $[0.1, 1.3]\bar{W}_t$; but the gap increases to 1.8 kids when people earn $[3.0, 4.1]\bar{W}_t$. Second, comparing the green with the orange line, we see vividly the combined price and income effect of social insurance on fertility. On the one hand, for relatively poor people, the price and income effects reinforce each other and push up the TFR. In Figure 2.1, for people whose income is within $[0.1, 1.1]\bar{W}_t$, TFR in the presence of social insurance is higher than that in the absence of social insurance. On the other hand, changes in fertility rate among relatively rich people depend on which of the two effects dominates. According to Figure 2.1, we observe a weak dominance of the income effect in that for parents earning $[1.2, 6.0]\bar{W}_t$, TFR with social insurance is lower than or equal to the one without social insurance. Third, the impact of the liquidity constraint is quite clear by comparing the orange with the purple line. If people can not borrow against their future pension benefits, their current period resources shrink after tax payments, which induces a lower fertility rate. This effect is particularly significant for the poor. For example, people earning $0.1\bar{W}_t$ choose not to have children if they can't borrow against their future pension benefits; but instead raise 1.3 kids if this liquidity constraint is absent. The magnitude of this effect diminishes when peoples' income rises. That is, if their income is sufficiently high, this constraint will not bind and thus have no impact on their fertility.

2.4.2 Average Saving Rate

According to our model, parents face two basic decisions: consumption and investment. The former one is divided into consumption of themselves and transfer to their elderly parents; while the latter one is divided into investment in children which is risky and in risk-free private savings, plus some mandatory social insurance contributions. Based on our previous discussion, the price effect of social insurance lowers the cost of childbearing, which induces an increase in the rate of return on child investment. Similarly, other factors such as expected long-run GDP growth and interest rate also affect the rewards on child investment. Since saving and fertility are close substitutes for parents in our model, we are interested to see whether saving behavior changes in an opposite way as compared to the fertility.

Figure 2.2 shows the average private saving rate as a function of parents' income. As expected, people tend to save more when their income rises. A comparison between the baseline scenario (the blue line) and the hypothetical scenario without social insurance in 2005 (the green line) indicates that the saving rate of the poor stagnates around 0.1%, while private saving among the rich increases significantly in response to the big drop in TFR as seen in Figure 2.1. What's more, there are two other notable observations when social insurance is in place. First, according to the orange line, relatively poor people whose income is less than $1.1\bar{W}_t$ have great incentives to borrow against their future pension benefits to raise children, reflected as the negative private saving rate at the current period. However, this borrowing incentive diminishes quickly when parents' income rises above the society average, indicating the loss of interest in borrowing as a

result of a decline in the rate of return on child investment for the rich. Second, when the liquidity constraint is applied, everyone's saving rate is positive, but the value is lower than the one without social insurance because people have fewer resources left after tax payments.

2.4.3 Transfer Rate

The third variable of potential interest is the transfer from the middle-age people to their elderly parents, which measures the cross-generational linkages between adult children and elderly parents.

Given our model specification, decisions on investment and consumption can be separated. Within consumption, because of logarithm utility, parents are maximizing a Cobb-Douglas aggregation of their own and their elderly parents' consumption²¹. Therefore, if there is no uncertainty in child investment, transfer and consumption each occupies a fixed proportion of the total expenditure. It is then quite natural to see a convergence in the transfer rate for people earning above average income, and some dynamics for people whose consumption decisions are affected by precautionary motive in the absence of social insurance and by pension benefit generosity in the presence of such an income redistribution program. From Figure 2.3, when parents' income is above the average, the transfer rate stabilizes around 20% without social insurance coverage; and around 23% with the coverage²².

²¹This is because

$$\log[C_t^m] + \beta \log[C_{t+1}^o] + \eta \log[C_t^o] = (1 + \beta + \eta) \log[(C_t^m)^{\frac{1}{1+\beta+\eta}} (C_{t+1}^o)^{\frac{\beta}{1+\beta+\eta}} (C_t^o)^{\frac{\eta}{1+\beta+\eta}}]$$

²²The transfer rate is defined as a percentage of disposable income. So a 23% transfer rate with social

From Figure 2.2, private saving rate in the presence of social insurance is always smaller than that without the program. With the mandatory social insurance saving, parents' burden to save enough resources for their own elderly life due to the precautionary saving motive lessens so that they can transfer more resources to their elderly parents, which helps explain the higher convergence level in transfer when social insurance is in place. In addition, although some factors consistently affect the rate of return on child investment and thus on fertility, they have almost no impact on transfer behavior when comparing the 1985 baseline with the hypothetical scenarios. This is another confirmation that fertility and private saving are close substitutes.

2.5 Sensitivity Analysis

Broadly speaking, the main factor that affects fertility rates is the rate of return on child investment, which could be decomposed into the childbearing cost ($a + bW_t$) and the expected transfer from children (D_{t+1}). First, since we do not have data to measure the fixed and income-varying costs of raising children, parents' degree of altruism towards grandparents, and the distribution of children's transfers, we make somewhat arbitrary assumptions on related parameters and suppose that their values did not change between 1985 and 2005. However, we conjecture that there might exist an increasing trend of childbearing cost in reality and the cost itself may be heterogeneous across income groups. Therefore, we provide sensitivity analysis on both the fixed and income-varying costs to show their impacts on fertility choices.

Second, based on our discussion of the channels through which social insurance can

insurance is roughly equal to 21%(= 23% × (1 - 8%)) of pre-tax income.

affect fertility, we consider the price and income effects as well as the liquidity constraint channel explicitly in our calibration exercises and heterogeneous analysis. However, in terms of potential changes in social attitudes that the degree of altruism may decrease and parents may expect a lower children's transfer rate in the presence of social insurance, it is hard to tell whether social attitude changes have already taken place, and if so, to measure the magnitude of the changes. In order to shed lights on the potential impact of this channel on fertility, we simulate our models under hypothetical scenarios that either η or \bar{d}_{t+1} ($\mathbb{E}_t[d_{t+1}]$) changes its value in 2005.

Third, social insurance program is not stationary over time and due to the progress of an aging society, it is expected that the social insurance tax rates will rise to raise more funds from the working generation in order to support the retired generation. The changes in social insurance tax rates will alter both the opportunity cost of childbearing and the expected life-time income. Boldrin et al. (2005) show that more than 50% of the differences in TFR between US and European countries can be explained by the differences in social insurance programs. They also find a significantly negative relationship between TFR and the social insurance tax rates. Thus, we are interested to see whether this relationship holds in China and whether Chinese parents are sensitive to the tax rate changes.

2.5.1 Childbearing Cost

In our benchmark calibration, we assume the fixed cost of raising a child equals one year's average income (i.e. $a = 0.05$). Here, we consider two candidate values of a to examine its effect on fertility. The simulated TFRs across parents' income distribution

are displayed in Figure 2.4.

The blue line represents the 1985 benchmark scenario and the green line shows the results in the presence of social insurance in 2005. The orange and the purple lines are simulated TFRs under two hypothetical scenarios when the fixed cost of childbearing increases to two year's or four year's average income (i.e. $a = 0.1$, or 0.2). It is clear that TFR is negatively correlated with the fixed cost and poor peoples' fertility rates are more sensitive to changes in the fixed cost. For example, for people earning $0.1\bar{W}_t$, they choose to have 1.3 kids when the fixed cost equals one year's average income; however, this number decreases to 0.6 if the fixed cost doubles; and it further drops to 0 if the fixed cost rises to four years' average income. By contrast, rich peoples' fertility rates are not affected much by this kind of change; the magnitude in TFR decline is less than 0.5 when people earn more than 3 times the average.

Regarding the income-varying cost, we treat it as 10% of parents' income W_t in the baseline scenario, and then increase it to 15% and 20%. We also find a negative relationship between simulated TFRs and the income-varying cost in Figure 2.5, but this negative relationship is different from the one between TFR and the fixed cost in several aspects. First, relatively poor people are not so sensitive to the changes in b . For people earning $[0.1, 0.5]\bar{W}_t$, the difference in TFRs is less than 1.0 when b increases from 10% to 20%. The poorest people still choose to have 0.9 kids, while we notice that they will choose not to have children when a rises from one year's to four year's average income. Second, changes in TFR for relatively rich people are more pronounced. For those whose incomes are within the range of $[0.6, 2.7]\bar{W}_t$, the decline in TFR is larger than 1.0. Third, we find that when couples' income is higher than $3.2\bar{W}_t$, they choose not to have children

if the income-varying cost rises to 20%. This result provides a supporting evidence for the low fertility rate among the rich: the high income-varying cost prohibits some of them from having children.

2.5.2 Potential Changes in Social Attitudes

The potential changes in social attitudes are twofold: one is parents' degree of altruism towards their elderly parents (η) may decrease; the other is parents expect that the distribution of their children's transfer rate may move leftwards. First, we consider two candidate values for η ($\eta = 0.35$, or 0.20) to examine how peoples' fertility rates change if they do not pay as much attention to their parents as before. Second, based on our assumption that transfer rate d_{t+1} follows a Beta distribution (i.e. $d_{t+1} \sim \text{Beta}(\alpha_d, \beta_d)$), we modify the value for β_d so that the average transfer rate declines from the original 12.5% to 10% or 8%.

The simulated TFRs under various η values are shown in Figure 2.6. For people earning $[0.1, 0.7]\bar{W}_t$, when η declines, TFR increases by a small amount. However, when parents' income is higher than $0.7\bar{W}_t$, there is almost no change in TFR when η drops from 0.5 to 0.2, which is a little bit surprising.

In order to explain this observation, we need to look at the changes in optimal transfer rate d_t when η decreases. Since η represents how grandparents' utility is viewed by parents, a smaller η means the same amount of transfer brings less rewards (i.e. utility) to parents²³. Figure 2.7(a) illustrates the interactions between η and d_t that when η declines

²³This can be proven in eq. (2.10), from which we can derive that $\frac{\partial d_t}{\partial \eta} > 0$. This shows a positive relationship between parents' altruism degree η and their optimal transfer rate d_t .

from 0.5 to 0.2, everyone's transfer declines and thus the transfer rate converges to a much lower level at around 10% compared to the previous 23%. At the same time, from Figure 2.7(b), when η falls, the private saving rate does not change for those whose incomes are within $[0.1, 0.7]\bar{W}_t$, but rises significantly for people earning $[0.8, 6.0]\bar{W}_t$. Combining Figure 2.7(a) and 2.7(b), a decline in η directly reduces the transfers from the middle-age people to their elderly parents, while how they allocate the extra resources depends on the comparison in the rate of return between risky child investment and risk-free private saving. For relatively poor people, rate of return on fertility is higher than that on saving so that we observe an increase in TFR; by contrast, relatively rich people see higher return on saving than on fertility, which induces an increase in s_t instead.

In contrast to the minor change in TFR when η declines, TFR decreases significantly when the mean expected old-age transfer from children declines. In Figure 2.8, when \bar{d}_{t+1} falls from 12.5% to 8%, the simulated TFRs decline for everyone, although the magnitudes are different depending on parents' income position. For people earning $[0.1, 0.5]\bar{W}_t$ or $[2.8, 6.0]\bar{W}_t$, the simulated TFR only changes a small amount (less than 0.5) even if \bar{d}_{t+1} drops by nearly 40%. However, people earning $[0.6, 2.7]\bar{W}_t$ are quite sensitive to changes in \bar{d}_{t+1} . For example, for people earning average income, their TFR drops from 2.4 to 1.15, when \bar{d}_{t+1} decreases from 12.5% to 8%.

2.5.3 Social Insurance Tax Rate

Due to the rapid development of an aging society, the aged dependency ratio rises dramatically in China. In a "pay-as-you-go" system, a challenging question is how to

ensure that this program is sustainable in the long run²⁴. There are two options: one is to increase the social insurance tax rates; the other is to cut benefits. In light of this, we simulate our model under hypothetical assumptions of higher social insurance tax rates and examine peoples' reactions in fertility.

As expected, an increase in social insurance tax rates can affect fertility through two channels. First, an increased tax rate α_t indicates a smaller childbearing cost ($b(1 - \alpha_t)W_t$) for everyone; according to the price effect, this may induce a higher fertility rate.

Second, the income effect of a rise in α_t is slightly complicated. Our specification on social insurance benefits consists of two parts: one is called “minimum pension” which is a small percentage (γ_t) of the average income level \bar{W}_t ; the other varies with workers' contribution. One important parameter is β_t , which can be viewed as “marginal actuarial fairness index”—the marginal pension increase for one more dollar's tax payment. To fund the minimum pension, it is necessary for the administration to impose a β_t less than 1. Such design of social insurance implies that poor people will get a better deal than the rich: the government-provided social insurance is an income redistribution from the rich to the poor. The net benefit from social insurance could be written as:

$$\text{Net Benefit at } t = \frac{\text{Benefit at } t + 1}{R_{t+1}} - \text{Contribution at } t \quad (2.16)$$

$$= \gamma_t \bar{W}_t + \beta_t \alpha_t (1 - bf_t) W_t - \alpha_t (1 - bf_t) W_t \quad (2.17)$$

$$= \gamma_t \bar{W}_t - (1 - \beta_t) \alpha_t (1 - bf_t) W_t \quad (2.18)$$

²⁴In our partial equilibrium framework, we do not explicitly specify a dynamic budget constraint for the government. In practice, with overlapping generations, there are arguments for a PAYG pension such as efficiency gains and intergenerational redistribution. However, there are tensions when the ratio of the number of contributors to the number of retirees decreases, which justify a need to raise taxes.

A threshold level of income (W_t^*), where the contribution to and the benefit from the program are equal, can be derived as follows:

$$W_t^* = \frac{\gamma_t \bar{W}_t}{(1 - \beta_t) \alpha_t (1 - b f_t)} \quad (2.19)$$

For families with income higher than the threshold level, the program reduces their present discounted life-time income; and vice versa for people with income below the threshold level. *Given f_t fixed*, eq. (2.18) implies a negative relationship between the net benefit and α_t . Hence a rise in α_t will reduce the net benefit for everyone (both the rich and the poor), and lead to a smaller W_t^* . Therefore we expect to see a lower f_t since children are normal goods. However this is not the end of the story: *given α_t fixed*, eq. (2.18) implies a positive relationship between the net benefit and f_t . An initial decline in f_t due to an increase in α_t will further reduce the net benefit, and thus leads to another decline in f_t .

Figure 2.9 shows the decomposition of the interactions between α_t and f_t . The purple line in 2.9(a) refers to the contribution as a function of income, and the blue one the benefit function. Hence, the gap between them constitutes the net benefit function as in 2.9(b), which is positive for $W_t < W_t^*$ and negative for $W_t > W_t^*$. When the social insurance tax rate increases from 8% to 12%, with f_t fixed, we see that the slopes of both the purple and the blue lines increase, but due to $\beta_t < 1$, the slope increase in the benefit function is smaller. This translates into a downward movement of the net benefit function, as displayed by the black dashed line in Figure 2.9(b). This represents a decline in the lifetime income for everyone (except for those with $W_t = 0$), and the direct income effect says people should demand fewer kids.

There is a further reaction in the model: with α_t now fixed at its new value, a smaller f_t induces higher slopes for both the contribution and the benefit functions, which in turn leads to an even lower net benefit function, as reflected in the black thin line in Figure 2.9(b). The second round of income effect will further reduce fertility. In the end, families respond to a higher social insurance tax rate by having fewer kids. The economy is characterized by high taxation and generous pension benefit (see the purple and blue thin lines in 2.9(a)), which resembles what we observe in those European welfare states.

It is interesting to note the difference in the income effect between introducing social insurance into the economy and raising tax rates within an existing system. In the former case, income effect is positive for the poor and negative for the rich. However, in the latter case, the income effect from a rise in α_t is negative for everyone earning a positive income. Combining this with the positive price effect, we must turn to calibration exercise to see how f_t responds.

We simulate our model under hypothetical scenarios that the current social insurance tax rate increases from 8% to 12% or 16%. Figure 2.10 shows that the magnitude of TFR reduction with respect to an increase in α_t is trivial, indicating an offsetting between the price and income effect.

Finally, we briefly discuss the effect of an increase in expected future tax rate α_{t+1} . This expectation change may happen if couples, for example baby boomers, anticipate a decline in the ratio of workers to beneficiaries next period. An increased future tax rate means a lower rate of return in child investment from the view of parents. Thus, parents will choose to have fewer kids, which reinforces their initial expectation of a smaller population size and a higher tax rate α_{t+1} .

2.6 Conclusion

This paper quantified the effect of social insurance on total fertility rate (TFR) through macro aggregation based on solving individual dynamic fertility choice model. From parents' point of view investment in children is risky and the risks are twofold: (1) whether a kid can grow up, represented by the infant mortality rate; (2) uncertainty about the transfer parents expect to get when they become old.

We examined the effect of social insurance on fertility rate from four perspectives. First, we considered the price effect due to payment of social insurance taxes, which decreases the opportunity cost of childbearing and thereby induces higher fertility. Second, we considered the income effect that this income-redistribution program has, which raises the expected life-time income for relatively poor people, while decreasing life-time income of the rich. Since children are viewed as normal goods, there is a positive relationship between life-time income and fertility. Combining these two effects, we find that low-earning people tend to have more kids since the two effects work in the same direction, while changes among high-income people depend on which of these effects dominates. Our heterogeneous analysis indicates that the income effect dominates for relatively rich people so that their fertility rate declines in the presence of the program.

The third factor we considered is the liquidity constraint caused by the need to pay the social insurance taxes, which reduces current cash-on-hand for everyone. Although it is possible to borrow through informal channels, this liquidity constraint plays an important role in affecting fertility choices. Comparing scenarios with and without this constraint, we find that social insurance causes average TFR to drop significantly. Specif-

ically, the fertility rate among the lowest earners where the taxes are most likely to cause binding liquidity constraint experiences the biggest decline.

The final aspect we took into account is the potential change in social attitudes, which is reflected in a decrease in either parents' altruism degree towards their elderly parents or expected transfers from their children. On the one hand, these changes in social attitudes were not addressed in previous literature. On the other hand, it is quite hard to examine whether this kind of change already happened after the social insurance program was introduced, and even harder to measure the magnitude of this change. In light of this, we used sensitivity analysis and simulated our model based on various candidate parameter values. Our results show that fertility rates are quite sensitive to changes in expectation of children's transfer rate; a 20% decrease will induce 10-70% drop in fertility rate depending on parents' position in the income distribution. In addition, a smaller degree of altruism induces less transfer from adult children to elderly parents in a direct and significant way. However, comparing the rate of return on private saving and child investment, people have heterogeneous choices in allocating the extra resources on either fertility or saving.

In contrast to Boldrin et al. (2005), who showed that at least 50% of the differences in TFR between US and European countries can be explained by the differences in the generosity of social insurance, we have shown that a more generous social insurance program has a trivial effect on the fertility rate in China. Beyond this, our calibration results match quite well with the declining tendency in TFR from 1985 to 2005. Decomposing the calibrated TFRs reveals that the combination of price and income effect of social insurance has a slightly positive impact on TFR, while imposing liquidity constraints could

decrease fertility rate by around 35%. Moreover, the decomposition brings out an important feature that factors related with the rate of return on child investment and private saving always matter when people are making fertility choices. Any changes in these factors have a direct effect on fertility rate, although the magnitude may vary across scenarios. Our calibrations show that, without consideration of liquidity constraint, the big drop in TFR from 1985 to 2005 is, to a large extent, due to changes in expected GDP growth rate and interest rates, which affect the average income in the children's generation and thus the future retirement income for their parents. In a similar spirit, we show that increases in childbearing costs also have significant effect on reducing fertility rate, although low-income people are more sensitive to changes in the fixed cost while high-income people are more sensitive to changes in the income-varying part.

There are many extensions to this work that could provide further insights on the effect of social insurance on fertility. First, our model is calibrated to match TFR in China between 1985 and 2005. Besides China, we observe a world-wide decreasing trend in TFR from the 1960s. It is interesting and important we calibrate the models to match other countries' TFRs during the same period. By doing this, we could have a better understanding of whether the same framework can be used among other countries, whether the effect of increasing tax rate on fertility is still limited, and whether those factors related to the rate of return on child investment also play an important role in other countries. Cross-country comparisons could help answer all these questions and provide a more comprehensive explanations on how the social insurance program affects peoples' fertility choices. Second, although our calibrations match actual TFR reasonably well, we make arbitrary assumptions about some key parameter values such as the cost

for raising children, the expected transfer rate, and parents' degree of altruism. It would be better if we could utilize micro-level data and empirically estimate these parameters by matching moments from simulations to the data. I leave these extensions for future work.

Table 2.1: Benchmark Parameters

Parameter	Symbol	Value
<i>Macro Economy</i>		
Number of years per period	T	20
Risk-free real interest rate 1966-1985	R_{1985}	1.889
Risk-free real interest rate 1986-2005	R_{2005}	0.940
Risk-free real interest rate 2006-2025	R_{2025}	1.220
GDP growth rate 1966-1985	g_{1985}	2.067
GDP growth rate 1986-2005	g_{2005}	5.213
GDP growth rate 2006-2025	g_{2025}	2.207
Average labor income at period t	\bar{W}_t	1.000
<i>Preferences</i>		
Discount factor	δ	0.818
Parents altruism towards grandparents	η	0.500
<i>Child Investment</i>		
Mortality rate between age 0 and 19	p	0.050
Childbearing cost: fixed part	a	0.050
Childbearing cost: income-varying part	b	0.100
Beta Distribution for transfer rate	α	2.000
	β	14.000
<i>Social Insurance</i>		
Marginal tax rate	α_t	0.080
Actuarial fairness index	β_t	0.800
Minimal pension ratio	γ_t	0.100
<i>State Variables</i>		
Grandparents' total fertility rate in 1965	f_{1965}	6.080
Grandparents' total fertility rate in 1985	f_{1985}	2.600
Grandparents' private saving rate in 1965	s_{1965}	0.050
Grandparents' private saving rate in 1985	s_{1985}	0.150

Table 2.2: Relative Income Distribution for Calibration

Percentile	0-5%	6-10%	11-20%	21-40%	41-60%	61-80%	81-90%	91-100%
Proposed w_{1965}	0.4629	0.5528	0.6480	0.7844	0.9500	1.1563	1.4077	1.9196
w_{1985}	0.4629	0.5528	0.6480	0.7844	0.9500	1.1563	1.4077	1.9196
w_{2005}	0.2414	0.3553	0.4595	0.6340	0.8733	1.2010	1.6507	2.7593
Means of w_{2025}	0.2414	0.3553	0.4595	0.6340	0.8733	1.2010	1.6507	2.7593

Table 2.3: Calibrated TFR and Private Saving Rate under Hypothetical Environments

		SINo	SIYes	SIYes+LC	Price & Income Effect	Liquidity Constraint Effect
1985	TFR	2.55	2.82	2.10	0.27	-0.72
	s_t	0.56%	-16.06%	0.65%	-16.62%	16.71%
2005	TFR	1.80	1.96	1.27	0.17	-0.69
	s_t	8.43%	-8.90%	6.99%	-17.33%	15.89%

Table 2.4: Effects of Other Factors on TFR

	TFR	s_t	d_t
Benchmark	2.55	0.56%	19.76%
R_{t+1}	2.49	1.42%	19.76%
g_{t+1}	2.21	4.85%	19.96%
g_t	2.55	0.56%	19.96%
f_{t-1}	2.55	0.57%	19.24%
$R_t s_{t-1}$	2.55	0.56%	19.57%
w_t distribution	2.37	0.85%	19.87%




Figure 2.1: Optimal Fertility Choice across Parents' Income




Figure 2.2: Optimal Private Saving Rate across Parents' Income




Figure 2.3: Optimal Transfer Rate across Parents' Income

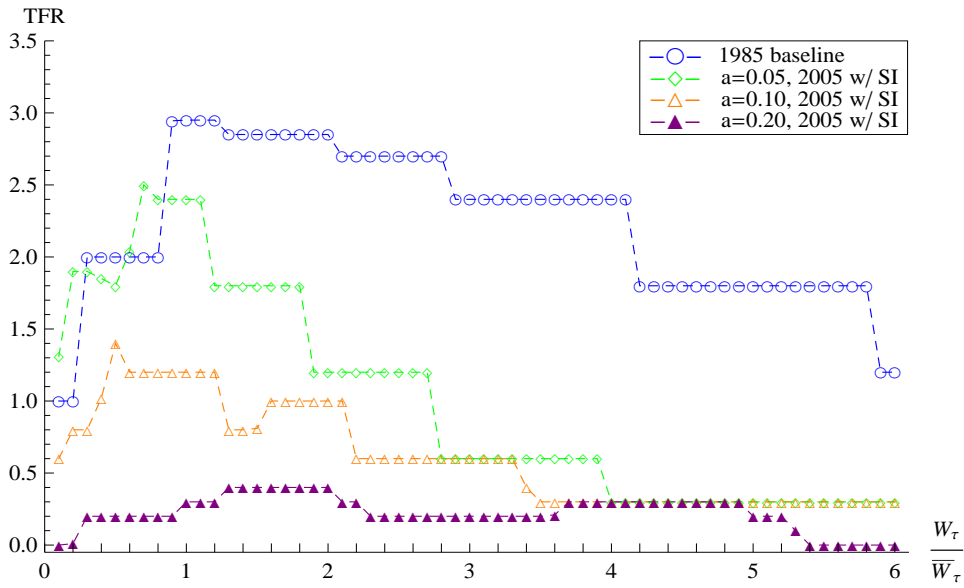


Figure 2.4: Optimal Fertility Choice and Fixed Cost of Childbearing a




Figure 2.5: Optimal Fertility Choice and Income-varying Cost of Childbearing b




Figure 2.6: Optimal Fertility Choice and Parents' Altruism Degree towards Grandparents η

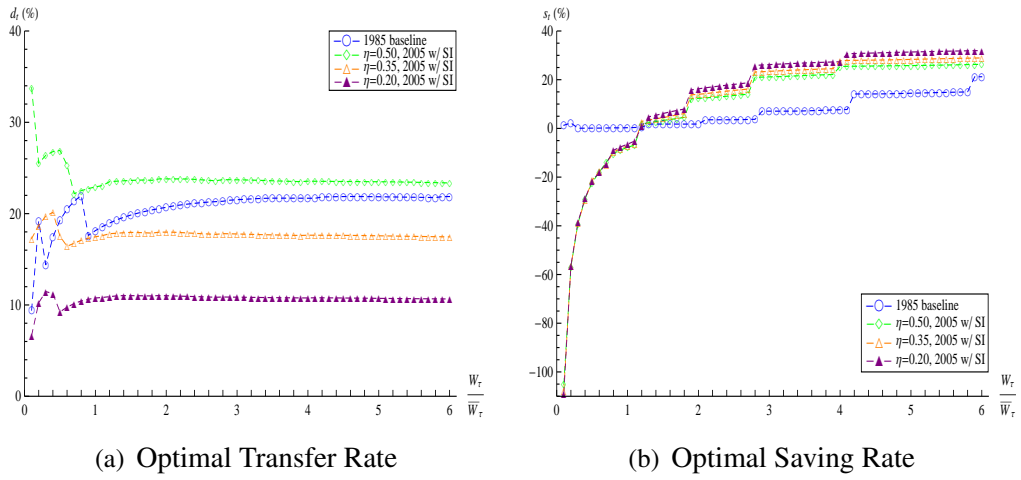


Figure 2.7: Optimal Transfer/Saving Rate and Parents' Altruism Degree towards Grandparents η






Figure 2.8: Optimal Fertility Choice and Expected Child Transfer Rate $\mathbb{E}_t[d_{t+1}]$

Contribution & Benefit



(a) Contribution and Benefit

Net Benefit



(b) Net Benefit

Figure 2.9: Income Effect of an Increase in α_t




Figure 2.10: Optimal Fertility Choice and Social Insurance Tax Rate α_t

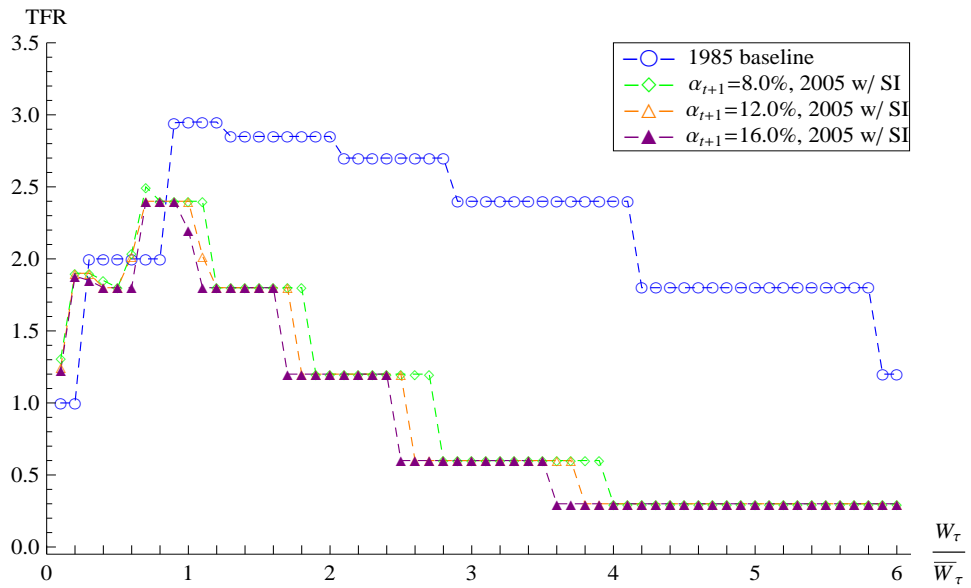


Figure 2.11: Optimal Fertility Choice and Expected Social Insurance Tax Rate α_{t+1}

Bibliography

- ANDERSON, S. (2003): "Why Dowry Payments Declined with Modernization in Europe but Are Rising in India," *Journal of Political Economy*, 111, 269–310.
- (2007): "The Economics of Dowry and Brideprice," *Journal of Economic Perspectives*, 21, 151–174.
- BARRO, R. J. AND G. S. BECKER (1989): "Fertility Choice in a Model of Economic Growth," *Econometrica*, 57, 481–501.
- BECKER, G. S. AND R. J. BARRO (1988): "A Reformulation of the Economic Theory of Fertility," *Quarterly Journal of Economics*, 103, 1–25.
- BHATTACHARJYA, D., A. SUDARSHAN, S. TULJAPURKAR, R. SHACHTER, AND M. FELDMAN (2008): "How Can Economic Schemes Curtail the Increasing Sex Ratio at Birth in China?" *Demographic Research*, 19, 1831–1850.
- BILLARI, F. C. AND V. GALASSO (2008): "What Explains Fertility? Evidence from Italian Pension Reforms," *IGIER Working Papers* 343.
- BOLDRIN, M. AND L. E. JONES (2002): "Mortality, Fertility and Saving in a Malthusian Economy," *Review of Economic Dynamics*, 5, 775–814.
- BOLDRIN, M., M. D. NARDI, AND L. E. JONES (2005): "Fertility and Social Security," *NBER Working Paper Series*, No. 11146.
- CALDWELL, J. C. (1978): "A Theory of Fertility: from High Plateau to Destabilization," *Population and Development Review*, 4, 553–577.
- (1982): *Theory of Fertility Decline*, New York: Academic Press, 6 ed.
- CARROLL, C. D. (2011): "Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems," *Working Paper, Department of Economics, Johns Hopkins University*.
- DAS GUPTA, M., W. CHUNG, AND S. LI (2009): "Is There an Incipient Turnaround in Asia's "Missing Girls" Phenomenon?" *World Bank Policy Research Working Paper Series*, No. 4846.
- DAS GUPTA, M., Z. JIANG, B. LI, Z. XIE, W. CHUNG, AND B. HWA-OK (2003): "Why is Son Preference so Persistent in East and South Asia? A Cross-country Study of China, India, and the Republic of Korea," *Journal of Development Studies*, 40, 153–187.
- EBENSTEIN, A. Y. (2011): "Estimating a Dynamic Model of Sex Selection in China," *Demography*, 48, 783–811.

- EBENSTEIN, A. Y. AND S. LEUNG (2010): "Son Preference and Access to Social Insurance: Evidence from China's Rural Pension Program," *Population and Development Review*, 36, 47–70.
- ECHEVARRIA, C. AND A. MERLO (1999): "Gender Differences in Education in a Dynamic Household Bargaining Model," *International Economic Review*, 40, 265–286.
- ECONOMIST (2010): "The War on Baby Girls: Gendercide," *The Economist*.
- EDLUND, L., H. LI, J. YI, AND J. ZHANG (2007): "More Men, More Crime: Evidence from China's One-Child Policy," *Institute for the Study of Labor (IZA) Discussion Papers*, No. 3214.
- EHRlich, I. AND J. KIM (2007): "Social Security and Demographic Trends: Theory and Evidence from the International Experience," *Review of Economic Dynamics*, 55–77.
- JUSTER, F. T. AND F. P. STAFFORD (1991): "The Allocation of Time: Empirical Findings, Behavioral Models, and Problems of Measurement," *Journal of Economic Literature*, 29, 471–522.
- LI, H., J. YI, AND J. ZHANG (2010): "Estimating the Effect of the One-Child Policy on Sex Ratio Imbalance in China: Identification Based on the Difference-in-Differences," *Institute for the Study of Labor (IZA) Discussion Papers*, No. 5149.
- LI, H. AND H. ZHENG (2009): "Ultrasonography and Sex Ratios in China," *Asian Economic Policy Review*, 4, 121–137.
- MA, Y., L. FENG, Y. CHEN, AND M. LENG (1998): *New Theory of Sex Ratio at Birth and Its Applications*, Capital University of Economics and Business Press (in Chinese).
- MODIGLIANI, F. AND S. L. CAO (2004): "The Chinese Saving Puzzle and the Life-Cycle Hypothesis," *Journal of Economic Literature*, 42, 145–170.
- QIAN, N. (2008): "Missing Women and the Price of Tea in China: The Effect of Sex-Specific Earning on Sex Imbalance," *Quarterly Journal of Economics*, 123, 1251–1285.
- RAO, V. (1993): "Dowry 'Inflation' in Rural India: A Statistical Investigation," *Population Studies*, 47, 283–293.
- RUST, J. (2008): "Dynamic Programming," in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, London: Palgrave Macmillan, Ltd.
- SANTOS, M. S. AND J. RUST (2003): "Convergence Properties of Policy Iteration," *SIAM Journal on Control and Optimization*, 42, 2094–2115.
- SEN, A. (1990): "More Than 100 Million Women Are Missing," *New York Review of Books*, 37.

- WEI, S. AND X. ZHANG (2011): “The Competitive Saving Motive: Evidence from Rising Sex Ratios and Savings Rates in China,” *Journal of Political Economy*, 119, 511–564.
- WILLIAMSON, N. (1976): *Sons or Daughters: A Cross-Cultural Survey of Parental Preferences*, Beverley Hills, CA: Sage Publications, 6 ed.
- WILLIS, R. J. (1982): “The Direction of Intergenerational Transfers and Demographic Transition: The Caldwell Hypothesis Reexamined,” *Population and Development Review*, 8, 207–234.
- ZHANG, J. AND J. ZHANG (2004): “How Does Social Security Affect Economic Growth? Evidence from Cross-Country Data,” *Journal of Population Economics*, 17, 473–500.