Consumption at old age and life time labor supply in rural China*

Yufeng Huang and Tobias J. Klein†

October 2013

Abstract

In many western societies, a large fraction of the consumption expenditures at old age is financed by a pension system. In rural China, as in many other developing countries, such a pension system does not exist. Therefore, consumption at old age is financed by savings and by transfers from children. The latter, however, is much more likely to be the case for sons, because daughters typically leave their family at the point in time at which they get married and join the families of their husbands. This partly explains why Chinese parents have a strong preference for having at least one boy who will then support them at old age. In this paper, using unusually rich data on time use and family relations we characterize the impact of the gender composition of the children on labor supply and retirement behavior for families in rural China. We then estimate a life cycle model to quantify the effect of the one child policy on well-being of individuals. Finally, we use the model to simulate the consequences of introducing a pension system for son preference. Our findings are not only relevant in the context of China because variation in the gender composition of the children generates a shock that is akin to a large shock to financial wealth—an event that may occur also in western societies, for instance due to a financial crisis. Studying the effects of such a wealth shock in the context of China allows us to study its effects separately from the impact of changes in labor market conditions. This is not possible for events like the financial crisis because they affect both, labor market conditions and wealth.

JEL Classification: J22, J26, J16.

Keywords: Retirement planning, labor supply, son preference, rural China, pension system.

*We would like to thanks Arthur van Soest for helpful comments.

†Huang: CentER, Tilburg University; Klein: CentER, TILEC, Tilburg University; E-Mail for correspondence: T.J.Klein@uvt.nl.

1 Introduction

In many western societies, a large fraction of the consumption expenditures at old age is financed by the pension system. In the context of a life cycle model, this means that individuals have to save less, compared to a situation without a pension system. Instead, they save by paying pension contributions.

In rural China, as in many developing countries, such a pension system does not exist. Therefore, consumption at old age is financed by savings and by transfers from children. The latter, however, is much more likely to be the case for sons, because daughters typically leave their family at the point in time at which they get married and join the families of their husbands. This partly explains why Chinese parents have a strong preference for having at least one boy who will then support them at old age. Transfers in this context may also consist of help received for performing daily activities such as cleaning the house, shopping or cooking. In this paper, we show that the gender composition of children affects labor supply and retirement behavior of their parents. This finding is not only relevant in the context of China. But it is particularly promising to study it in the context of China because on the one hand the introduction of the one child policy and changes therein create institutional variation that make it possible to identify the effects of interest, and on the other hand we can make use of unusually rich unusually rich data on time use and family relations.

The one child policy was introduced in 1979 (see e.g. Hesketh, Lu, and Xing, 2005, for more details). The objective was to slow down population growth. It had a negative effect on those rural families who had to rely on having at least one boy. This is because only a boy could continue the family business, as it was not easy to sell it at a fair price (such market failures often arise in developing economies). For that reason selective abortions of girls became especially common in rural areas (Kim, 2005; Ebenstein, 2011; Leung, 2011). The policy was subsequently adapted so that families in rural China could have a second child if the first was a daughter.

In this paper, we characterize the impact of the gender composition of the children on labor supply during work age and retirement behavior for families in rural China. We then estimate a life cycle model to evaluate the effect of the one child policy on well-being of individuals through its impact on retirement and savings behavior and labor supply. Finally, we use the model to simulate the consequences of introducing a pension system for son preference.

Our research question is how gender composition of children affects labor supply and retirement choices made by their parents. In terms of the economics, having a daughter instead of a son is akin to a negative wealth shock early in life, in the form of the change in expectation that one's pension benefits,

including transfers from children, will be lower than previously expected. In this sense, our paper does not only have relevance for China, but for many developing countries in which rapid growth often leads to such wealth shocks—in one way or another. Studying this question in the context of China is promising because the introduction of the one child policy and changes therein create institutional variation that makes it possible to identify the effects of interest. At the same time, the income effect of having a girl rather than a boy is similar to a very large negative wealth shock well before retirement—something that could come about in western societies as well, for instance when people expect that pension benefits will be less generous or when they are hit by a financial crisis. Therefore, studying the effect of this on consumption choice and labor supply is not only interesting in the context of developing countries, but also more generally within the field of the economics of aging.

Next to addressing this more general question we contribute to the scarce literature on labor supply and retirement behavior in China. China will become a more and more important player in the world, and therefore it is important to extend our understanding of aging-related questions to China. It is well documented that there is selective abortion, but it is to our knowledge not documented how this influences labor supply and retirement behavior. In that sense, our goal to provide clean evidence on this is already a major step forward.

Our paper is related to the literature on consumption choice. Meghir and Pistaferri (2011) survey recent developments in the literature that studies how the dynamics of earnings and wages affect consumption choices over the life cycle. Gourinchas and Parker (2002) estimates a structural model of optimal life-cycle consumption expenditures in the presence of realistic labor income uncertainty. Banks, Blundell, and Tanner (1998) show that consumption drops at retirement. Laitner and Silverman (2005), Skinner (2007) and Hurst (2008), among others, argue that this may be explained by non-separabilities between consumption and leisure.

It is also related to the literature on (female) labor supply. See Killingsworth and Heckman (1986), Heckman (1993), Blundell and MaCurdy (1999), and Blundell, Macurdy, and Meghir (2007) for excellent surveys of this literature. One strand of the literature focuses on static considerations—the tradeoff between consumption and leisure in a given period (see e.g. van Soest, 1995). Another strand of the literature focuses on dynamic considerations (see e.g. Heckman and MaCurdy, 1980; French, 2005). Here, the decision is between retiring now, which gives more leisure but less consumption over the life cycle, or later. We contribute to both strands by proposing a model in which the gender of the child gives rise to an income effect and thereby affects labor supply over the entire life cycle. Adda, Dustmann, and Stevens (2011) analyze the life-cycle career costs associated with child rearing and decomposes their effects into

unearned wages (as women drop out of the labor market), loss of human capital, and selection into more child-friendly occupations.

The decision to supply labor at old age is closely related to retirement choice and its effects. Rust and Phelan (1997) analyze how the U.S. Social Security and Medicare insurance system affects the labor supply of older males in the presence of incomplete markets for loans, annuities, and health insurance. French (2005) investigates the effects of health, wealth, and wages on labour supply. Rendon (2006) examines the relationship between wealth accumulation and job search dynamics. He proposes a model in which risk-averse individuals search for jobs, save, and borrow to smooth their consumption. French and Jones (2011) analyze the effects of employer-provided health insurance, Medicare, and Social Security on retirement behavior. Gustman and Steinmeier (2005) find that there is wide heterogeneity in time preferences among individuals. Those with high discount rates often retire at 62. They have few assets and heavily value lost benefits from working after 62, largely ignoring potential increases in later benefits. Declining actuarial adjustments beginning at 65 induce those with low discount rates to retire at 65. Van der Klaauw and Wolpin (2008) estimate a model of retirement and savings incorporating limited borrowing, stochastic wage offers, health status and survival, social security benefits, Medicare and employer-provided health insurance coverage, and intentional bequests. Gustman and Steinmeier (2000) estimate a structural econometric model of retirement of dual-career couples using panel data from the National Longitudinal Survey of Mature Women. The estimates suggest that there is a correlation of tastes for leisure between the two spouses, which generates in a correlation of labor supply.

The Chinese economy, and thereby the labor market has seen dramatic changes over the last decades. Therefore, our paper is related to the following two papers on the relationship between economic development and individual choices. Eckstein and Lifshitz (2011) estimate a dynamic model of female labor supply and find that the rise in education levels accounts for about 33 percent of the increase in female employment, and the rise in wages and narrowing of the gender wage gap account for another 20 percent, while about 40 percent remains unexplained by observed household characteristics. They show that this unexplained portion can be empirically attributed to cohort-specific changes in preferences or the costs of child-rearing and household maintenance. Attanasio, Low, and Sánchez-Marcos (2008) construct a lifecycle model of female participation and savings, and investigate which changes in the determinants of labor supply account for the increases in participation early in the life-cycle observed for the 1950s cohort. They find that a combination of a reduction in the cost of children alongside a reduction in the wage-gender gap is needed.

The value to having a boy is related to geographic proximity. Compton and Pollak (2011) show that

close geographical proximity to mothers or mothers-in-law has a substantial positive effect on the labor supply of married women with young children. They argue that the mechanism through which proximity increases labor supply is the availability of childcare.

Next, in Section 3, we provide details on the institutions in rural China. The data set we use is described in Section 2. In our model, parents value having a boy because he will ultimately support them and will most likely stay close to them. The main reduced-form evidence is presented in Section 5. Thereafter, in Section 6, a life cycle model is developed and estimated that allows us to simulate the effects of the introduction of a pension system. This cannot be done without having such a model in hand, because generations who are affected are only about to retire. Finally, Section 7 summarizes and concludes.

2 The China Health and Nutrition Survey

Our data come from the China Health and Nutrition Survey (CHNS) database—an on-going longitudinal survey conducted by the Carolina Population Center. It is primarily designed to answer policy-relevant questions regarding health, nutrition and family planning policies. Yet, its broad scope of questions enables studies on labor supply decisions.

The surveyed households are drawn from 9 provinces from the eastern part of China.[1] In order for the sample to be representative, within each province, the surveyors stratify its counties by income and select 4 of the counties that represent the income distribution of the province. Among the 4 counties, the provincial capital is almost always selected (with 2 exceptions when another large city was selected instead of the capital), together with one of the lower income cities. Within a county, neighborhoods (in urban/suburban areas) and villages (in rural areas) are drawn randomly, and are considered to be the primary sampling units. There are in total 190 primary sampling units before 1993, and the number is later increased, up to 216 by 2000.[2] Within a primary sampling unit, approximately 20 households are drawn and followed over time, though the exact number varies due to attrition.

The initial wave of the survey was conducted in 1989, when a total of 3,795 households were sampled, within each, all pre-schoolers and adults aged between 20 and 45 were surveyed. From the second wave, which was conducted in 1991, all individuals were surveyed regardless of age. From the 1993 survey, new households formed by individuals who are from the previously surveyed households are added

[1] The 9 provinces are Guangxi, Guizhou, Heilongjiang, Henan, Hubei, Hunan, Jiagnsu, Liaoning, and Shandong. See http://www.cpc.unc.edu/projects/china/proj_desc/chinamap for a map.

[2] Before 1993 these were 32 urban neighborhoods, 30 suburban neighborhoods, 32 towns (county capital city), and 96 rural villages. After 1993, this changed to 36 urban neighborhoods, 36 suburban neighborhoods, 36 towns, and 108 villages.

to the sample, together with other new households to replace the dropped out ones. Up to the most recent 2009 survey, there are 6,910 households ever participated in the survey.

Individuals are identified using both, household and individual identifiers. Therefore, it is in principle straightforward to link household members. This brings information on joint labor supply and income. Because the percentage of widowed and divorced females is very low, and the reason for missing income or labor supply information for one of the spouses is unclear, we only use observations where information on both parties is observable.

For couples (or females) and their children, we restrict ourselves to women who have been married in at least one of the waves of the survey, and were younger than age 52 in her earliest participation in the survey. For this sub-sample, each respondent provides information on all the children she has given birth to, for their identity given by the household and individual identifier, gender, year and month of birth and whether they are alive. We ignore all cases of abortion or infant mortality, because these reported numbers are rare and, if provided, unreliable.

Finally, the respondents sometimes give information on their parents, in particular whether they are alive and the location of their residence. We link information on parents to the household data, without knowing the identity of the parents.

Among all the survey questions, we are most interested in household demographics, labor supply and fertility histories. The demographic and labor force participation surveys are conducted with each adult members of the household. Income data is imputed from summation of income from labor markets, receipts from non-market activities (for example, revenue earned from household agricultural produce), and revenues from assets and government subsidies. For each female participant who is ever married (i.e. married, divorced or widowed), the survey asks for the full pregnancy and child-bearing history (date of childbearing, gender and whether the child survives), from which we infer the child composition in the household. Next, we describe the key variables in detail.

The primary interest is in the labor supply of the couple (female and male), in terms of hours of work per week. In the adult survey questionnaires, every adult was asked the labor supply in, and earnings from a variety of sources, including contracted jobs, household productions and other paid non-contracted jobs, as well as income earned from land and other assets. These categories include primary and secondary (formal) occupation, home gardening, collective and household farming and fishing, raising livestocks, small handicraft or commercial business, as well as retirement and other forms of government subsidies.

For each of these categories, the raw data contains the number of hours per typical working day, the

Table 1: Sample selection

	nr. females – urban b	nr. females – rural b	nr. obs – urban b	nr. obs – rural b
Full sample	2182	4283	9506	23148
– with info on children and household size	1968	3907	9215	22385
– and labor supply	1966	3897	7681	17676
– and income	1783	3616	6207	14953
– and husband labor/income	1403	3053	4757	12119

Notes: An observation is a unique couple-wave combination, with identity for both the female and the male respondent.

number of days per working month, and the number of months worked for the past year. We compute the average hours of work per week in the past year using these numbers. Income data is already aggregated into annual income by type of work.[3] We consider these data as the annual nominal income.

While most of the income data is nominal, the raw data reports both nominal and real total household income. The surveyor collected local price data in each area of the survey, and use it (together with the standard consumer basket supplied by the State Statistics Bureau) to calculate annual CPI per province, separately for rural and urban areas. Real income is computed using this CPI measure. While the CPI measure is not supplied by the data directly, we backward-compute the implied CPI from nominal and real total income, and use it to deflate all income measures to their 2008 price level in a given province-rural/urban area.

Our sample contains information on households over time for which we have complete information on birth records, labor supply, and income. Table 1 shows that this information is available for 3053 out of 4283 households in rural areas. While we focus on rural households, here and in the following we also provide information on urban households in order to highlight the similarities and differences between the two. The total number of observations over time is 12119, which means that on average we observe a household about four times.

3 Institutional environment

China has seen rapid economic growth over the last decades, in particular in urban areas. Ge and Yang (2012) use a national sample of Urban Household Surveys and document that between 1992 and 2007, the average real wage increased by 202 percent, accompanied by a sharp rise in wage inequality. They

[3]For example, non-retirement income is constructed by months worked times the monthly earnings, plus annual bonus and other cash income earned from each of the jobs that the individual participated in. Retirement income is the sum of the reported retirement pensions or salaries that the individual reports.

Figure 1: Age of wife at time of marriage

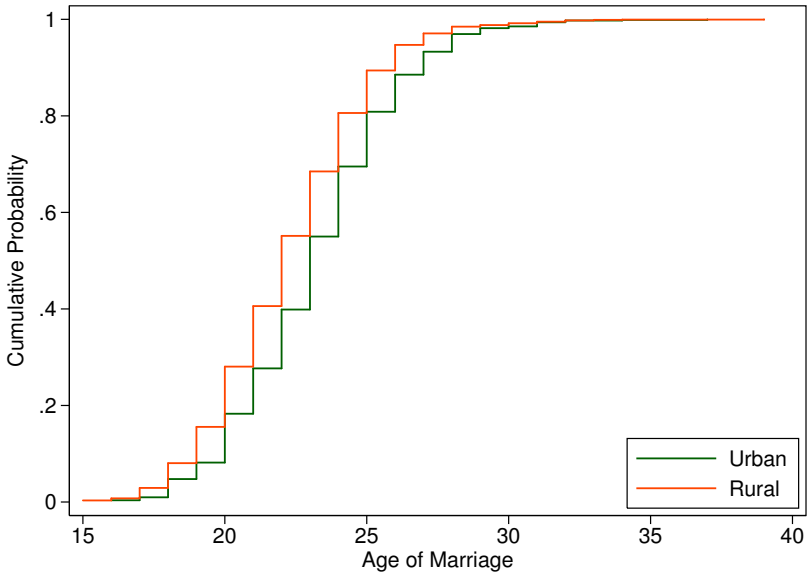

attribute this to higher pay for basic labor, rising returns to human capital, and increases in the state-sector wage premium.

In many ways, the social norms and the institutional environment in rural China favor conducting an analysis like ours because many choices that are made in western societies, and thereby complicate an analysis, are less of a choice in rural China. First, the date of marriage is not a choice in the way it is in western societies. Generally speaking, it is considered a norm that marriage takes place at relatively young age. Figure 1 confirms this. It shows that the age of the wife at the time of marriage if relatively low. About 80 percent of the women are married by the age of 25. We can see that there is no big difference between urban and rural China.

Second, also the time at which the first child is born is not a choice. The reason is that it is considered a norm that a first child is born as soon as possible after the wedding. Figure 2 shows that more than 60 percent of the couples have a first child after 2 years. Also here we can see that there is no big difference between urban and rural China.

Third, the Chinese *hukou* household registration system restricts location choices (see e.g. Wu and Treiman, 2004). Therefore, it is typically not an option to move from a rural area to a city.

Taken together, these three reasons speak in favor of the viewpoint that the couple's main choices are labor supply and consumption, something we will argue when we formulate the structural model in Section 6.

For educational achievement, or more generally the quality-quantity tradeoff for children, the evi-

Figure 2: Timing of first birth

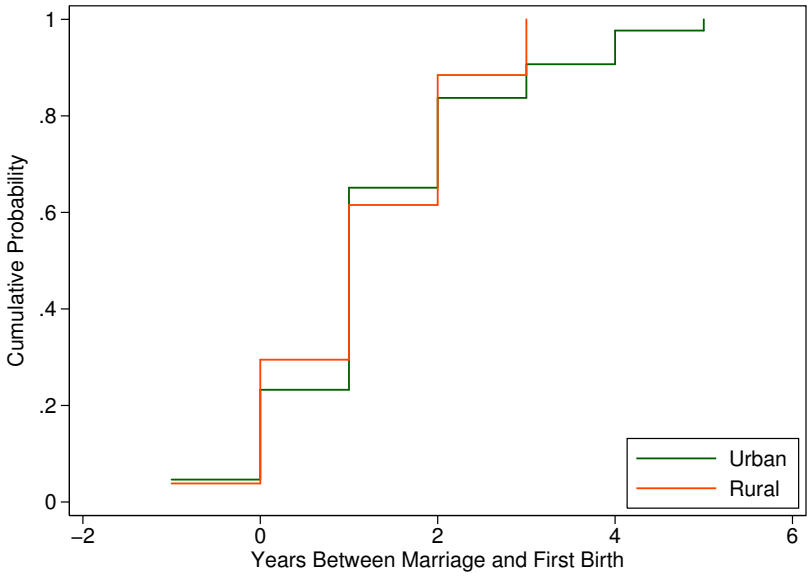

dence is not conclusive so far. Li, Zhang, and Zhu (2008) find that twins receive less education, while Qian (2009) finds that the increase in the number of children that is due to the relaxation of the one child policy over time led to more educational achievement.[4] In our data, we have not found strong evidence that speaks in favor of less educational achievement in bigger families. Rather, the patterns we observe seem to confirm the common perception that in China, the more education the better, no matter what and independent of the gender of the child.

An issue that has received great attention and that is closely related to our paper is the imbalance in the sex ratio. The sex ratio is defined as the number of newborn boys relative to the number of newborn girls and the natural ratio is 1.04. In China, however, this ratio has recently increased dramatically, in particular in rural China, as documented in Figure 3.

This has been explained by a pure preference for sons, together with the ability to perform selective abortion. Kim (2005) investigates the effect of prenatal sex selection on fertility through a stochastic dynamic model with uncertainty in conception as well as in gender, where a woman makes decisions on conception and abortion with or without gender detection tests (i.e. sex-selective or sex-unselective abortion). The paper shows that, when the cost of gender detection test falls, the sex ratio at birth rises due to more selective abortions, but fertility can rise or fall with rising sex ratio. Ebenstein (2011) presents

[4] Chang and Li (2012) examine the educational effects of sibship sex composition using data from the 2002 and 2007 Taiwan Social Change Survey. The results show that men with older siblings, especially older sisters, have higher educational attainment, while women with both younger sisters and younger brothers have lower educational attainment. Moreover, the positive effect of having an additional older sibling on men's educational attainment only exists in large families but not small families.

Figure 3: Sex ratio over time

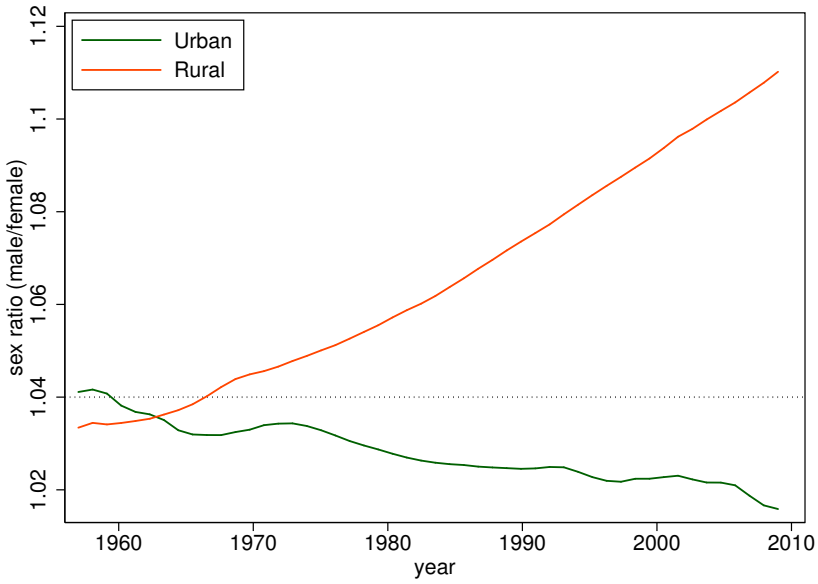

a model of fertility choice when parents have access to a sex-selection technology and face a mandated fertility limit. By exploiting variation in fines levied in China for unsanctioned births, he estimates the relative price of a son and daughter for mothers observed in China's census data from 1982 to 2000. He finds that a couple's first son is worth 1.42 years of income more than a first daughter, and the premium is highest among less-educated mothers and families engaged in agriculture.[5] Also Leung (2011) studies the effect of the one child policy on the Han-Chinese in urban area and the one-son-two-child policy in the rural area by studying the variations of sex ratios and fertility rates that is associated with different birth orders. We will pursue a similar Angrist, Lavy, and Schlosser (2010)-type identification strategy in our reduced-form analysis below. Almond, Edlund, and Milligan (2009) study immigrants to Canada and argue that the evidence they find for selective abortion speaks in favor of a preference for sons. The preference for sons, irrespective of where it originates, has important consequences. Wei and Zhang (2011) argue that as the sex ratio rises, Chinese parents with a son raise their savings in a competitive manner in order to improve their son's relative attractiveness for marriage and that this pressure on savings spills over to other households.

In this paper, we argue that the preference for sons does not only originate in cultural norms, but is also related to economic reasons. In particular, being close to one's children does not only allow parents to enjoy spending time with them—seeing children as a consumption good as in Wolpin (1987)—, but also allows them to draw on for help in the household or by providing informal care. In that sense,

[5] Akbulut-Yuksel and Rosenblum (2012) provide evidence on the consequences of the expansion of prenatal ultrasound use on sex-selection in India.

Figure 4: Relative distance to parents and to parents-in-law (rural female)

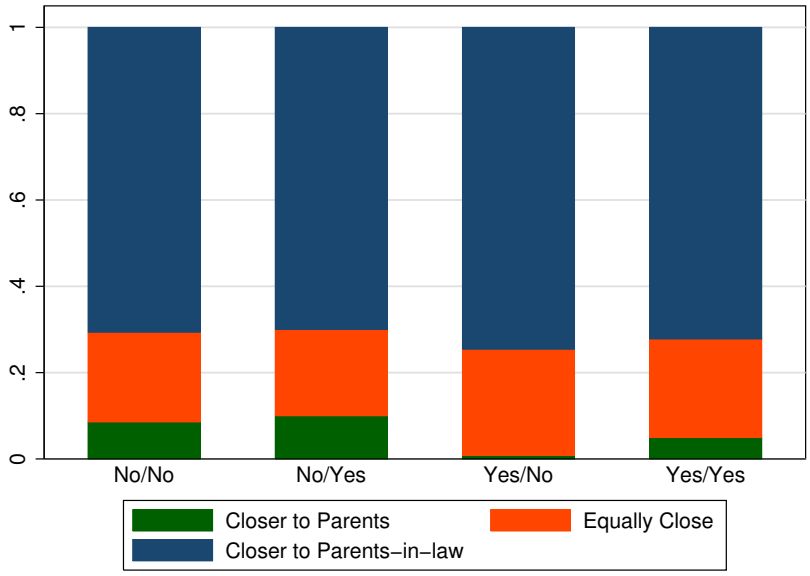

Notes: The horizontal axis denotes "whether the female has any brothers; whether her husband has any brothers"; so "No/Yes" means the husband has brother(s) while the wife does not. The graph is conditional on the female is married, and at least one of the parents as well as at least one of the parents-in-law are alive.

children can also be seen as investment goods. However, importantly, the social norms in China mean that usually, the wife will join the family of her husband, which also means that the couple will locate closer to the parents of the husband.

Figure 4 shows that the vast majority of couples lives closer to the parents of the husband. Importantly, this is the case irrespective of the gender composition of the siblings of both the husband and the wife. At the same time, children who live closer to their parents are more likely to receive a bequest (Yin, 2010).

4 Fertility and labor supply

Table 2 provides summary statistics, separately for households in rural and urban areas. It is interesting to make this distinction because it highlights the differences between those two groups. The most important difference in the context of this paper is that households in rural areas are bigger on average, because couples have more children. Total real income is lower in rural areas, and especially income from formal employment for women, something we discuss in more detail in Section 5.

It is interesting to also look at differences across cohorts. Table 3 provides averages over all observations where the wife was between 35 and 45 years old, by birth year of husband and wife, respectively.

Table 2: Summary statistics

	urban households		rural households	
	mean	std.	mean	std.
age of the female	37.68	11.38	38.46	12.51
number of household members	3.96	1.35	4.32	1.54
total mumber of children given birth to	1.65	0.83	2.25	1.09
total real income (female)	7715.11	11035.61	5611.88	10754.33
real income from formal jobs (female)	4663.97	9487.19	1571.54	5883.43
total hours per week (female)	35.83	31.68	33.95	31.94
hours per week on formal job (female)	20.67	22.94	8.99	19.07
total real income (male)	10374.65	15789.51	7869.70	16314.43
real income from formal jobs (male)	7653.71	14948.95	3785.45	12402.75
total hours per week (male)	39.43	31.76	39.86	33.20
hours per week on formal job (male)	24.87	24.05	14.61	22.40
obs.	4757	4757	12119	12119

Notes: Number of household members excludes members who moved out and reside in another household (even in the same neighborhood). All real income measures are normalized to 2008 CNY using the CPI at the province level. Formal jobs are defined as occupations with formal contracts, which includes contracted farm work, but excludes household farming activities. Total number of children ever given birth to excludes abortion and mortality at childbirth.

First, averages per couple were taken, and then the average across couples was calculated. For the number of children, the maximal number of children per couple was computed instead of the average, before averaging across couples. Recall in that respect that our data are for the years 1989 until 2009. The number of observations is provided in the last column and is for couples.[6] The last row provides statistics across all birth cohorts.

We see that the number of children has decreased dramatically and that in rural areas, any cohort had more children on average as the corresponding cohort in urban areas. At the same time, we see the well-documented pattern that the less children a family had on average over time, the more boys were born relative to girls. Younger generations work more and earn higher incomes.

[6] It is lower for the youngest cohort because of the requirement that only observations were used when the wife was between 35 and 45 years old. This was never met by the youngest individuals in our sample, because those wives born in 1979 were not yet 35 years old in 2009.

Table 3: Comparison between rural and urban households by cohort

rural

birth year	num. children	ratio boys to girls	income wife	hours wife	income husband	ho husband	obs.
1940-1949	2.986	1.034	3377.859	18.865	4478.345	31.827	585
1950-1959	2.213	1.104	4615.263	30.788	6129.566	35.204	1078
1960-1969	1.639	1.194	9547.486	39.342	13820.278	45.554	1306
1970-1979	1.335	1.254	13504.898	39.748	17021.266	49.721	898
total	2.044	1.146	7761.377	32.186	10362.364	40.576	3867

urban

birth year	num. children	ratio boys to girls	income wife	hours wife	income husband	ho husband	obs.
1940-1949	2.137	1.198	4846.925	23.815	5397.337	26.844	255
1950-1959	1.531	1.008	6121.065	33.990	7364.580	37.646	652
1960-1969	1.235	1.062	12566.483	34.579	15522.146	39.760	753
1970-1979	1.149	0.986	17027.676	39.141	22207.025	44.749	375
total	1.513	1.063	10140.537	32.881	12622.772	37.250	2035

Notes: This table is produced using data for individuals between age 35 and 45. The number of children and related indicators are computed by taking the maximum, while the income and hours measures are computed from the mean over all waves of surveys. The number of household members excludes members who moved out and reside in another household. All real income measures are normalized to 2008 CNY using the CPI at the province level. Formal jobs are defined as occupations with formal contracts. This includes contracted farm work, but excludes household farming activities. The total number of children ever given birth to excludes abortion and mortality at childbirth.

Figure 5: Time use of husbands and wives

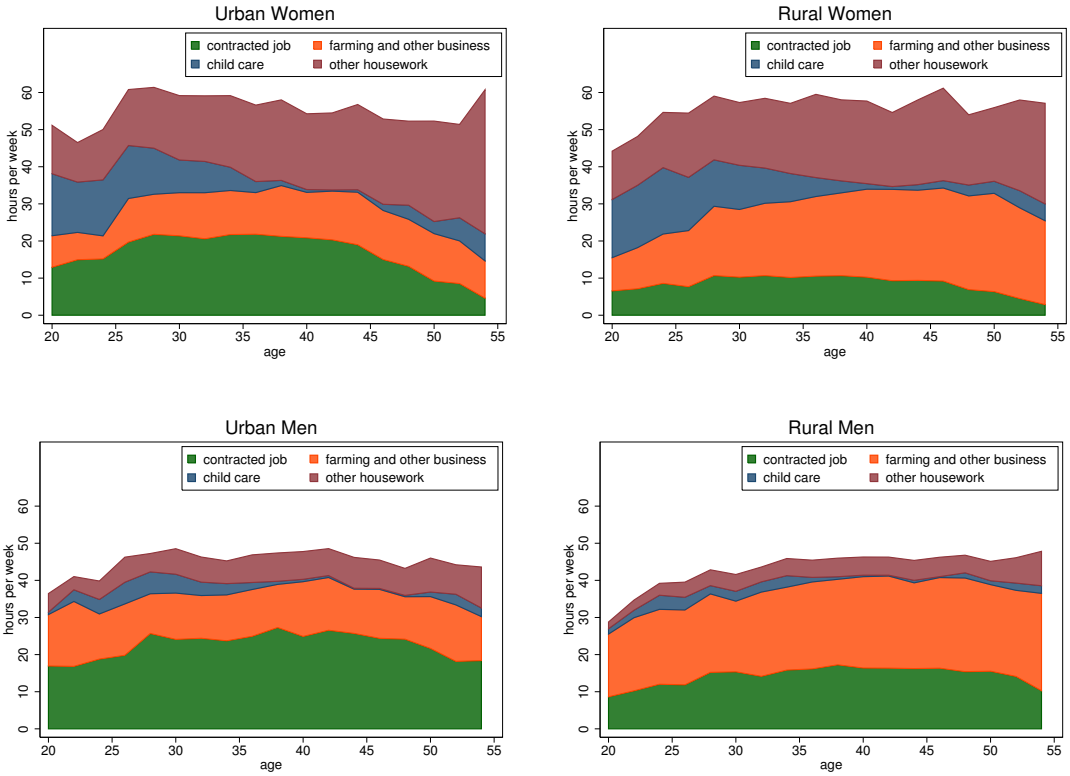

A nice feature of the CHNS is that it contains information of the time use by both, husbands and wives. This can be used to document the similarities and differences between urban and rural areas. Figure 5 shows that in both areas, women have their first child relatively early and spend a similar amount of time on child care until age 35. Also the amount of time spent on housework is similar in urban and rural areas. What is generally different is the importance of farming as a source of income in rural areas, relative to working in a formal, contracted job.

Figure 6 shows how different sources of income changed over time, separately for urban and rural areas, men and women, and individuals below and above the age of 50. These figures show that farming is by far the most important source of income for women in rural areas above the age of 50. Pension income has become more and more important over time, but foremost in urban areas. Moreover, income from a formal employment relationship has become more and more important also for young men in rural areas.

Figure 6: Income composition before and after retirement

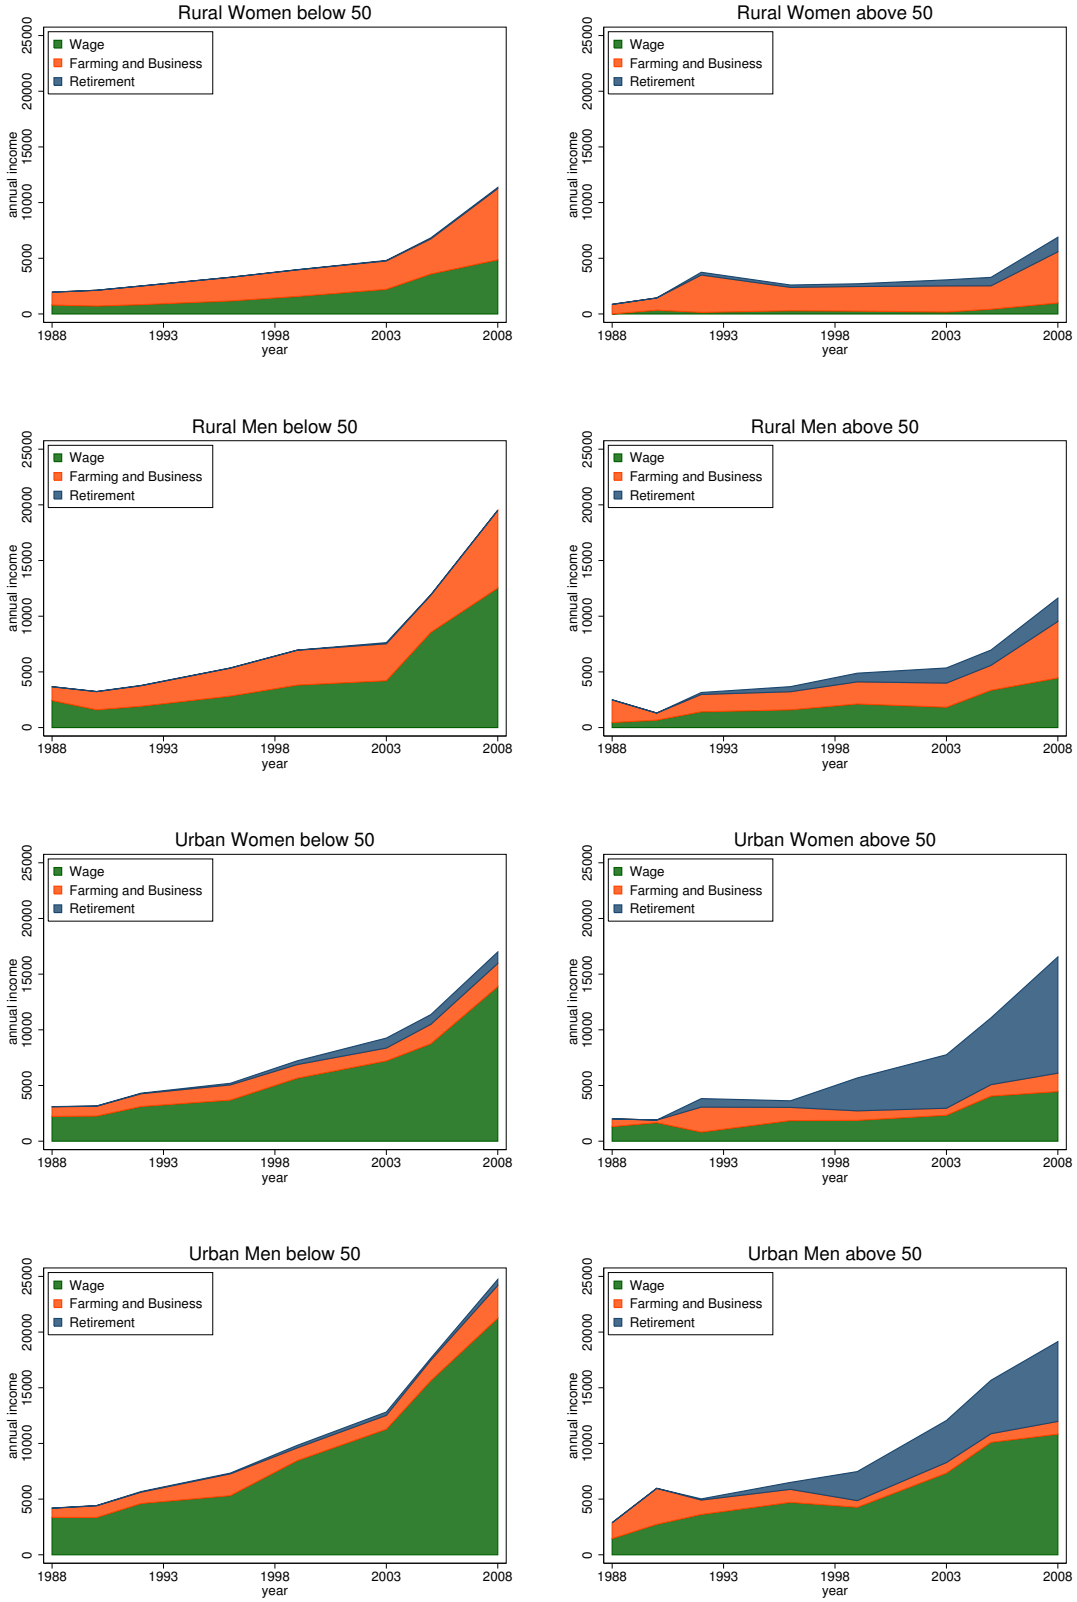

Notes: Income is real income in 2008 CNY.

Figure 7: Labor supply around childbirth

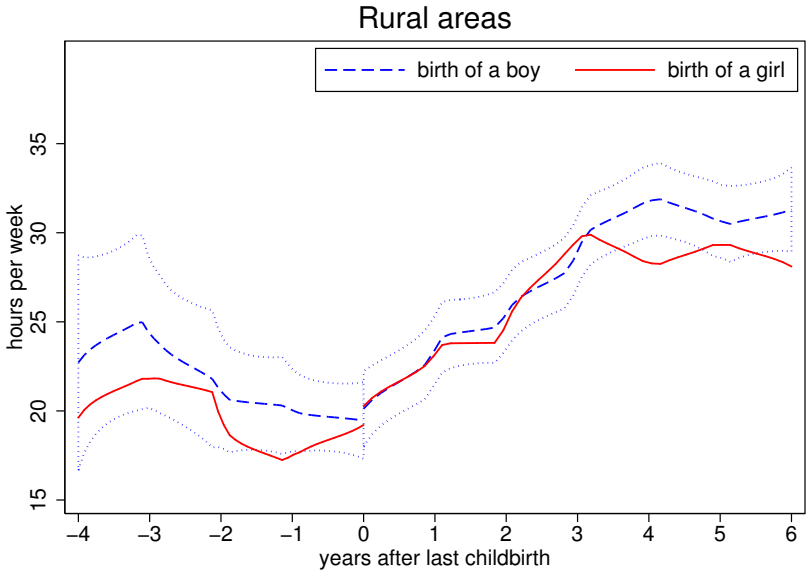

5 The effect of the gender composition of children on labor supply

The gender composition of the children in a family is not exogenous. This is because on the one hand, there is selective abortion and on the other hand, families who have no boy yet may opt for another child. This means that families who benefit more from having at least one boy will be more inclined to have a second or third child or opt for selective abortion, and therefore will be more likely to have at least one boy. For that reason, a regression of labor supply on gender composition and the number of children may yield biased results.

We address this concern by exploiting institutional changes that affected different generations of parents in different ways. In particular, we will exploit the timing of the introduction of the one child policy as well as the fact that rules were applied less strictly in later years. From this, we will be able to obtain clean estimates of the effects of the number of children and the gender composition on labor supply and income earned over the life cycle. We expect that he effects of having at least one boy will be that his parents retire earlier and work less over the life cycle. We will also be able to investigate whether his parents work harder at earlier stages of the life cycle to be able to invest more in a boy than they would in a girl (Wei and Zhang, 2011).

Before doing so, we shall point out that the focus of this paper is *not* the impact of fertility on labor supply during the time in which children grow up and receive education. In fact, due to the institutional and cultural environment, mothers in China go back to work relatively quickly after giving birth. To

Figure 8: Family lifetime labor supply and income

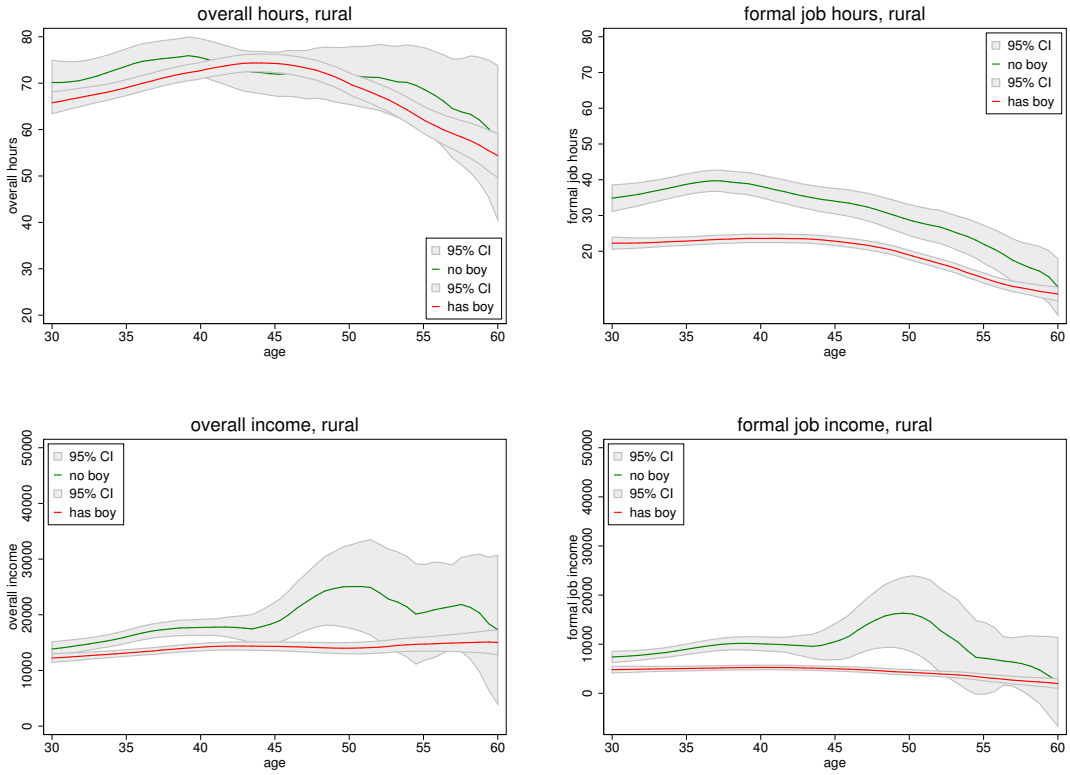

Notes: A family is defined as a couple. Income and hours are thus summed up across husband and wife and plotted against age of the wife.

provide descriptive evidence documenting this, we plot in Figure 7 the hours worked by the wife against years around the last childbirth.[7] We do so by gender of the child. Interestingly, the pattern for rural China resembles the one that was documented by Attanasio, Low, and Sánchez-Marcos (2008) for recent years in the U.S. The confidence interval is point-wise and for those instances in which a boy was born. The line for the case in which the child was a girl almost always lies within that confidence interval. This means that gender of the child does not seem to have an impact on labor supply around child birth.

The focus of this paper is on the relationship between lifetime labor supply and the gender composition of the children. Figure 8 plots hours worked, and income, by the husband and the wife, against the age of the wife. It distinguishes between families with a son and families without a son. In the context of rural China, overall hours worked are not a good measure of labor supply because they also include fairly mild work activity at home, with possibly some unrecorded idle time in-between. Therefore, we focus on hours in a formal job and income. The overall picture that arises is that labor supply is lower in

[7] The discussion in this section focuses on rural households. Evidence for urban households is provided in an appendix.

those families with a son.

Of course, for the reasons given before, the patterns documented so far could be entirely driven by confounding factors. To address this challenge, we exploit two sources of exogenous variation: the gender of the first child and institutional changes that affected different generations of parents in different ways. The former is a valid instrument if there is no selective abortion for the first child, which is an assumption that is testable because the "natural" ratio in the population is known. In our data, 52.3 percent of the first born children are male in rural areas, providing support of this assumption. At the same time, conditional on the first child being a girl, the probability that the second child is a boy is 1.39 times higher than the probability that a girl is born. This is the well-known number that is often related to selective abortion.

The institutional changes we exploit are related to the one child policy. In China, as documented in Section 3, most couples get a first child as soon as possible after getting married. Therefore, the date of marriage is indicative of the constraints the one child policy imposes on different couples. We divide couples into three groups and assume that after controlling for a flexible time trend the only difference between the three groups is in those constraints. The first group consists of couples who got married before 1975. The idea is that they were not affected by the one child policy. The second group consists of couples who got married between 1975 and 1985. These couples were the first to be constrained by the one child policy. After 1985, the rules were adjusted and interpreted in a more strict way, so couples who got married between 1985 and 2000 form the third group. After the year 2000, the regulations were changed again. Therefore, couples who got married after that year form the third group. Finally, we use as an instrument whether couples belong to an ethnic minority. The idea here is that these couples were exempted from the one child policy. Throughout, we control for cohort, year and province dummies.

There are two endogenous variables, the total number of children and whether at least one of the children is a boy. The respective first stage estimates are reported in the first two columns of Table 4. If a first child is female, the expected number of children is higher by 0.312 and the probability that none of the children will be a boy increases by 36.2 percentage points. Relative to the reference group of couples who got married between 1985 and 2000, couples who were not exposed to the one child policy have 0.569 children more on average and are less likely to have no boy at all, by 2.8 percentage points. The second group, which faced less strict constraints than the reference group, also has more children on average, 0.410, and is less likely to have no boy by 10.3 percentage points. Couples who got married after 2000 have 0.367 children less on average and are more likely to have no by, by 10.3 percentage points. Finally couples who belong to an ethnic minority have 0.131 children more on average and are

Table 4: Impact of gender on labor market outcomes for rural couples

	First: Number of Children	First: No Boy	Total Hours	Total Income	Formal Hours	Formal Income
45<age<60	0.105	0.294***	16.776**	2319.091	-1.322	918.085
	(0.169)	(0.063)	(7.148)	(2711.557)	(4.450)	(1999.259)
Gender first born	0.312***	0.362***				
	(0.018)	(0.007)				
marriedbefore75	0.569***	-0.028**				
	(0.034)	(0.013)				
marriedbefore85	0.410***	-0.103***				
	(0.037)	(0.014)				
marriedafter00	-0.367***	0.103***				
	(0.101)	(0.038)				
Ethnic Majority	0.131***	-0.042***				
	(0.029)	(0.011)				
numberchildren			-2.010	1338.007	-8.287***	-1434.655*
			(2.694)	(1021.885)	(1.677)	(753.446)
No Boy			3.823	-630.819	7.862***	1300.206
			(3.766)	(1428.673)	(2.344)	(1053.376)
45<age<60*No Boy			36.875***	13196.800***	18.255***	7630.046***
			(9.584)	(3635.506)	(5.966)	(2680.496)
age>60*No Boy			109.144***	6608.321	33.729	-502.749
			(37.091)	(14069.918)	(23.089)	(10373.896)
45<age<60*Nr. Child			-9.914***	-1965.280**	-2.226	-963.703
			(2.594)	(983.822)	(1.614)	(725.382)
age>60*Nr. Child			29.557***	-3283.164	16.550**	-1071.343
			(11.055)	(4193.505)	(6.882)	(3091.915)
cohort dummies	Yes	Yes	Yes	Yes	Yes	Yes
year dummies	Yes	Yes	Yes	Yes	Yes	Yes
province dummies	Yes	Yes	Yes	Yes	Yes	Yes
obs.	11970	11970	11970	11970	11970	11970

Notes: We select couples who i) reside in rural areas, for which ii) the income for both the wife and the husband are observed, and iii) where the wife is aged above 30. The number of children and whether there is a boy, and their interactions with age and age squared, are instrumented by gender of the first child, whether the household has an ethnic minority, whether the couple was married before 1985 or before 1975, or married after 2000, and their interactions with age and age squared.

slightly less likely to have no boy.

The outcomes we study are total hours, total income, formal hours and formal income earned by couples in which the wife is below the age of 45—our reference group, between 45 and 60 years old, and above 60 years old. Focusing on the number of formal hours, we find that the total number of children has a significant negative effect, at the one percent level—each additional child is estimated to cause a decrease in the number of formal hours by 8.287 when the wife is below the age of 45. It is not significantly different when the wife is between 45 and 60 years old, but above that age, more children actually go along with an increased number of hours. One explanation for this could be that then, the children have already left the house.

Importantly, these are the effects for couples who have at least one boy. Having no boy goes along with a significant additional increase in the number of formal hours by 7.862 when the wife is below the age of 45. Between 45 and 60 the number of formal hours is significantly higher, by 18.255, and after that it is even higher, but not significantly so. This shows that the gender composition has important effects on lifetime labor supply.

6 A structural life cycle model

We now formulate a structural model that will allow us to quantify the value of having at least one boy. This value comes from the transfers the couple receives from that boy at old age. For simplicity we will use a unitary model in which the household (and not the household members) has a utility function. Household utility will depend on consumption and leisure of both, husband and wife. In the model, the number of children and gender composition is not a choice. That is, the model is for consumption and labor supply—as well as retirement—from age 45 of the wife onwards, after investments into children have been made. The main difference between having a boy or not will be that a boy will provide his parents with extra utility when they are old, while a girl will leave the family. This extra utility comes from either monetary transfers, or helpingt them with their daily activities, or the like. This will generate an income effect and will therefore affect labor supply during working age, and retirement behavior, because it is a substitute for consumption that is financed by own wealth. We will estimate this model by indirect inference. Here, model parameters are chosen such that data generated by the model produce the same coefficient estimates in reduced-form regressions as those that come from the real data and are reported in Table 4. Thereby, as explained above, we exploit the institutional variation the introduction of the one child policy and changes therein created. In particular, we exploit the timing of the introduction

of the one child policy as well as the fact that rules were applied less strictly in later years. We then use the model to evaluate the effects of the one child policy on the welfare of families and to simulate the consequences of introducing a pension system on son preference.

6.1 Flow utility

We specify preferences to be of the iso-elastic-Cobb-Douglas type. Denote consumption in period t by c_t and leisure by l_t. Flow utility in t is

$$u_t = \frac{z_t^{1-\nu}}{1-\nu},$$

with

$$z_t = {c_t}^{\gamma} {l_t}^{1-\gamma}.$$

This specification is similar to the one in French and Jones (2011). In addition, couples receive utility Δ_b in each period above the age of 60 when they have at least one boy.

The quantity of leisure of the couple is

$$l_t = \bar{l} - n_t,$$

where $\bar{l}$ is the maximal possible leisure, n_t are working hours.[8]

6.2 Income process

Both spouses face the same wage process, with different observed initial conditions. We assume that the logarithm of wages at time t ($\ln w_{it}$), is a function of age (age_{it}), hours worked (h_{it}), and an autoregressive component (ω_{it}),

$$\ln w_{it} = wage(age_{it}, h_{it}, type_{it}) + \omega_{it}.$$

The autoregressive component has correlation coefficient ρ_w and is given by

$$\omega_{it} = \rho_w w_{it-1} + \eta_{it}, \quad \eta_{it} \sim \mathcal{N}(0, \sigma_\eta^2).$$

This is similar to the specification used by French (2005), Attanasio, Low, and Sánchez-Marcos (2008), and French and Jones (2011).[9]

[8] This is the sum of both spouses' leisure. We also consider a richer model in which leisure of the husband and the wife enter separately.

[9] Abowd and Card (1989) study individual earnings and hours data from three different longitudinal surveys and find that it is very similar across data sets, and may be adequately summarized by a simple components-of-variance model, consisting of

Following van Soest (1995) individuals' hours choice (h_{it}) is discrete. Income y_{it} consists of hours worked times wages,

$$y_{it} = w_{it} h_{it}.$$

6.3 Budget constraint

Assets in the following period are

$$a_{i,t+1} = a_{it} + y_{it} - c_{it}$$

and we impose the no-borrowing constraint

$$a_{it} \geq 0$$

for all t.

6.4 Decision problem

Individuals maximize the expected discounted sum of utilities,

$$U_t = \left(\sum_{\tau=t+1}^{T} \delta^{t-\tau} S_{t,\tau} u_\tau \right) + \delta^{T+1-\tau} S_{t,T+1} V_{T+1},$$

where δ is the discount factor, $S_{t,\tau}$ is the probability to be alive in τ from the viewpoint of period t and V_{T+1} is the continuation value when reaching period $T+1$.

It depends on the bequest function

$$b(a_{it}) = \theta_B \frac{(a_{it} + \kappa)^{(1-\nu)\gamma}}{1-\nu}.$$

This function is similar to the one used in French and Jones (2011).

This yields to the recursive formulation

$$V_t(a_t) = u(c_t, l_t) + \delta S_{t,t+1} \mathbb{E}\left[V_{t+1}(A_{t+1}) | c_t, l_t \right].$$

(i) serially uncorrelated measurement error, (ii) a shared component of earnings and hours with a second-order moving average covariance structure, and (iii) a nonstationary component that affects only the variances and contemporaneous covariances of earnings and hours. Guvenen (2007) studies the restrictions imposed by income processes on consumption data to distinguish between the two views. He finds that the life-cycle model with a process with individual-specific profiles and shocks with modest persistence, which has not been studied in the previous literature, is consistent with several features of consumption data, whereas the model with similar life cycle income profiles across individuals and very persistent shocks is consistent with some, but not with others. Meghir and Pistaferri (2004) study heterogeneity in the conditional variance of income shocks. They allow for education- and time-specific differences in the stochastic process for earnings and for measurement error. The conditional variance of the income shocks is modelled as a parsimonious ARCH process with both observable and unobserved heterogeneity. Using data from the 1967–1992 Panel Study of Income Dynamics they find strong evidence of sizable ARCH effects as well as evidence of unobserved heterogeneity in the variances.

We solve the model by backward recursion on a grid for wealth and interpolate between the knots when simulating expected values, where the expectation is over wage shocks ω_{it}.[10]

6.5 Estimation results

—to be added—

6.6 Policy simulation: the introduction of a pension system

—to be added—

7 Summary and concluding remarks

In this paper, we show that labor supply over the life cycle is generally related to the gender composition of children. In particular, we document that having no boys in rural China leads to an increase in labor supply and delayed retirement. Our proposed explanation is that daughters leave the family, while boys stay with their parents and provide for them at old age. This can be either in the form of monetary transfers, or by helping them with their daily activities.

Our findings are not only relevant in the context of China because more generally, within the realm of the life cycle model, variation in the gender composition of the children generates a wealth shock that is akin to a large shock to financial wealth—an event that may occur also in western societies, for instance due to a financial crisis. We have shown that studying this in the context of China allows us to disentangle the effects of such a wealth shock from the effects of changed labor market conditions. This is not possible for events like the financial crisis that affect both, labor market conditions and wealth.

[10] This is similar to Keane and Wolpin (1994) approach to use Monte Carlo integration to stimulate the required multiple integrals at a subset of the state points, and interpolating the non-simulated values using a regression function. Rust, Hall, Benıtez-Silva, and Pauletto (2005) compares different strategies for solving numerically dynamic programming problems that come from economic applications. Arcidiacono, Bayer, Bugni, and James (2012) introduce a method for approximating the value function of high-dimensional dynamic models based on sieves.

References

ABOWD, J., AND D. CARD (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica: Journal of the Econometric Society*, pp. 411–445.

ADDA, J., C. DUSTMANN, AND K. STEVENS (2011): "The Career Costs of Children," IZA Discussion Paper.

AKBULUT-YUKSEL, M., AND D. ROSENBLUM (2012): "The Indian ultrasound paradox," .

ALMOND, D., L. EDLUND, AND K. MILLIGAN (2009): "O sister, where art thou? The role of son preference and sex choice: evidence from immigrants to Canada," NBER Working Paper 15391.

ANGRIST, J., V. LAVY, AND A. SCHLOSSER (2010): "Multiple experiments for the causal link between the quantity and quality of children," *Journal of Labor Economics*, 28(4), 773–824.

ARCIDIACONO, P., P. BAYER, F. BUGNI, AND J. JAMES (2012): "Approximating high-dimensional dynamic models: Sieve value function iteration," .

ATTANASIO, O., H. LOW, AND V. SÁNCHEZ-MARCOS (2008): "Explaining changes in female labor supply in a life-cycle model," *The American Economic Review*, 98(4), 1517–1552.

BANKS, J., R. BLUNDELL, AND S. TANNER (1998): "Is there a retirement-savings puzzle?," *American Economic Review*, pp. 769–788.

BLUNDELL, R., AND T. MACURDY (1999): "Labor supply: A review of alternative approaches," *Handbook of labor economics*, 3, 1559–1695.

BLUNDELL, R., T. MACURDY, AND C. MEGHIR (2007): "Labor Supply Models: Unobserved Heterogeneity, Nonparticipation and Dynamics," *Handbook of Econometrics*, 6, 4667–4775.

CHANG, Y., AND J. LI (2012): "Budget Constraints and Son Preference in Educational Investment in Taiwan," .

COMPTON, J., AND R. POLLAK (2011): "Family Proximity, Childcare, and Women's Labor Force Attachment," NBER Working Paper 17678.

EBENSTEIN, A. (2011): "Estimating a dynamic model of sex selection in China," *Demography*, 48(2), 783–811.

ECKSTEIN, Z., AND O. LIFSHITZ (2011): "Dynamic female labor supply," *Econometrica*, 79(6), 1675–1726.

FRENCH, E. (2005): "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour," *Review of Economic Studies*, 72(2), 395–427.

FRENCH, E., AND J. JONES (2011): "The Effects of Health Insurance and Self-Insurance on Retirement Behavior," *Econometrica*, 79(3), 693–732.

GE, S., AND D. YANG (2012): "Changes in China's wage structure," .

GOURIEROUX, C., AND A. MONTFORT (1995): "Testing, Encompassing, and Simulating Dynamic Econometric Models," *Econometric Theory*, 11, 195–228.

GOURINCHAS, P.-O., AND J. A. PARKER (2002): "Consumption over the Life Cycle," *Econometrica*, 70(1), 47–89.

GUSTMAN, A., AND T. STEINMEIER (2000): "Retirement in dual-career families: a structural model," *Journal of Labor Economics*, 18(3), 503–545.

——— (2005): "The social security early entitlement age in a structural model of retirement and wealth," *Journal of Public Economics*, 89(2), 441–463.

GUVENEN, F. (2007): "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?," *The American Economic Review*, pp. 687–712.

HECKMAN, J. (1993): "What has been learned about labor supply in the past twenty years?," *The American Economic Review*, pp. 116–121.

HECKMAN, J., AND T. MACURDY (1980): "A life cycle model of female labour supply," *The Review of Economic Studies*, pp. 47–74.

HESKETH, T., L. LU, AND Z. W. XING (2005): "The Effect of China's One-Child Family Policy after 25 Years," *New England Journal of Medicine*, 353(11), 1171–1176, PMID: 16162890.

HURST, E. (2008): "The retirement of a consumption puzzle," NBER Working Paper 13789.

KEANE, M., AND K. WOLPIN (1994): "The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte Carlo evidence," *The Review of Economics and Statistics*, 76(4), 648–672.

KILLINGSWORTH, M., AND J. HECKMAN (1986): "Female labor supply: A survey," *Handbook of labor economics*, 1, 103–204.

KIM, J. (2005): "Sex selection and fertility in a dynamic model of conception and abortion," *Journal of Population Economics*, 18(1), 41–67.

LAITNER, J., AND D. SILVERMAN (2005): "Estimating life-cycle parameters from consumption behavior at retirement," NBER Working Paper 11163.

LEUNG, M. (2011): "A Structural Model of Sex Selective Abortion: The Effect of One-child Policy on Sex ratio imbalance in China," Mimeographed.

LI, H., J. ZHANG, AND Y. ZHU (2008): "The quantity-quality trade-off of children in a developing country: Identification using Chinese twins," *Demography*, 45(1), 223–243.

MEGHIR, C., AND L. PISTAFERRI (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72(1), 1–32.

MEGHIR, C., AND L. PISTAFERRI (2011): "Earnings, consumption and life cycle choices," *Handbook of Labor Economics*, 4, 773–854.

QIAN, N. (2009): "Quantity-quality and the one child policy: The only-child disadvantage in school enrollment in rural China," Discussion paper, National Bureau of Economic Research.

RENDON, S. (2006): "Job Search and Asset Accumulation under Borrowing Constraints," *International Economic Review*,

47(1), 233–263.

RUST, J., G. HALL, H. BENITEZ-SILVA, AND G. PAULETTO (2005): "A Comparison of Discrete and Parametric Approximation Methos for Solving Dynamic Programming Problems in Economics," Mimeographed.

RUST, J., AND C. PHELAN (1997): "How social security and medicare affect retirement behavior in a world of incomplete markets," *Econometrica: Journal of the Econometric Society*, pp. 781–831.

SKINNER, J. (2007): "Are You Sure You're Saving Enough for Retirement?," *Journal of Economic Perspectives*, 21(3), 59–80.

VAN DER KLAAUW, W., AND K. WOLPIN (2008): "Social security and the retirement and savings behavior of low-income households," *Journal of Econometrics*, 145(1), 21–42.

VAN SOEST, A. (1995): "Structural Models of Family Labor Supply: A Discrete Choice Approach," *Journal of Human Resources*, 30(1), 63–88.

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(3), 511–564.

WOLPIN, K. (1987): "Estimating a structural search model: the transition from school to work," *Econometrica: Journal of the Econometric Society*, pp. 801–817.

WU, X., AND D. TREIMAN (2004): "The household registration system and social stratification in China: 1955–1996," *Demography*, 41(2), 363–384.

YIN, T. (2010): "Parent–child co-residence and bequest motives in China," *China Economic Review*, 21(4), 521–531.

Tables and figures for urban households

Figure 9: Living arrangements for urban female

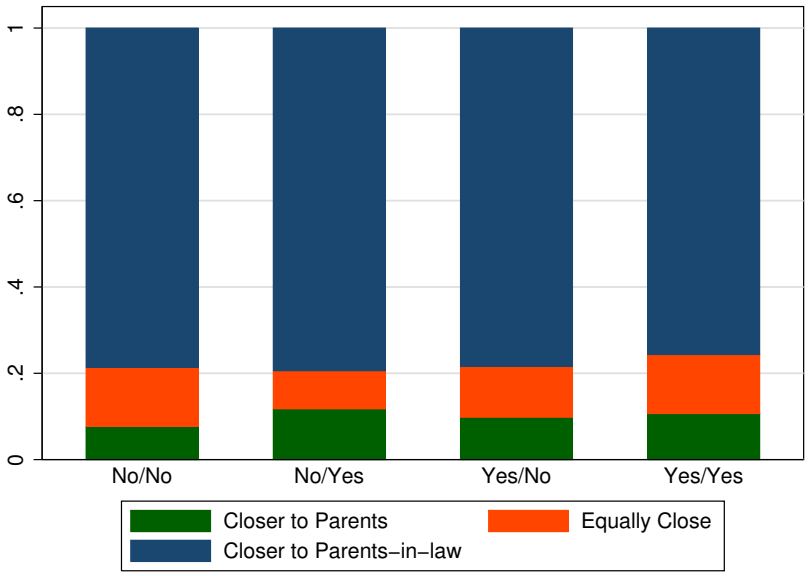

Notes: See figure 4.

Figure 10: Labor supply around childbirth

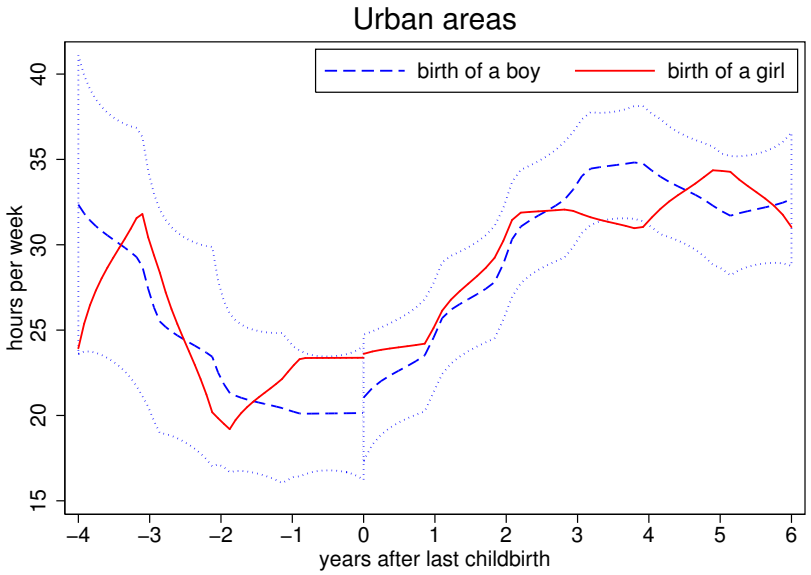

Figure 11: Labor market outcomes and gender of children for urban couples

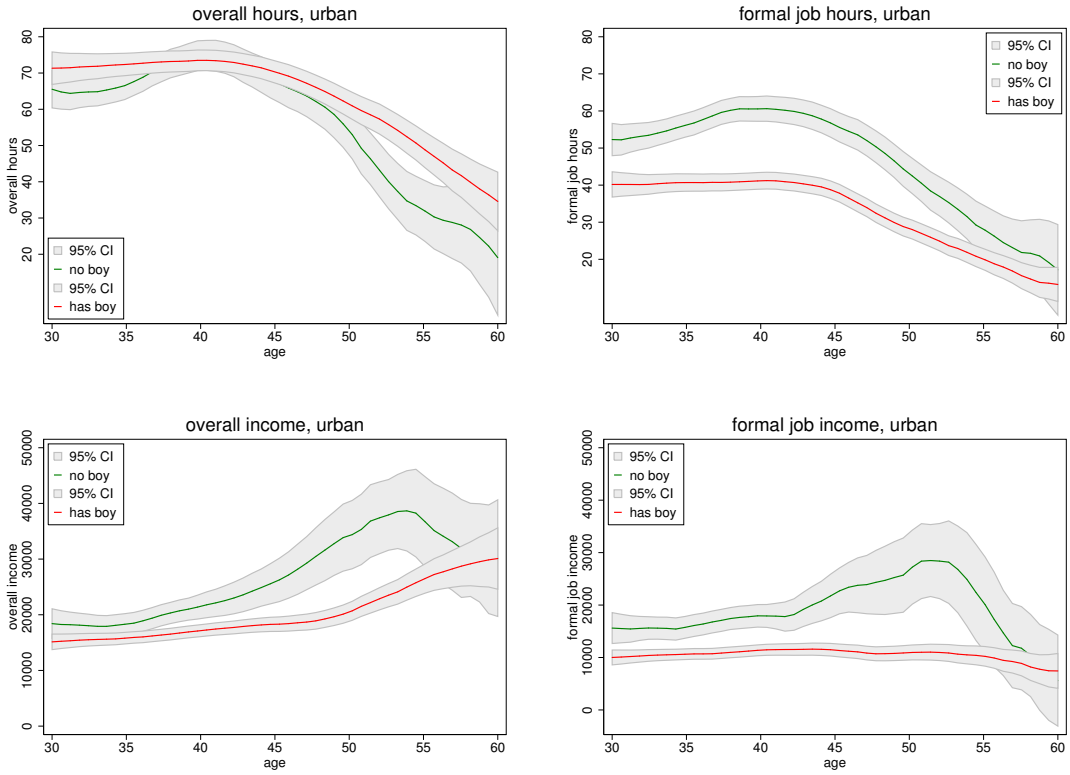

Notes: A family is defined as a couple. Income and hours are thus summed up across husband and wife and plotted against age of the wife.

Table 5: Impact of gender on labor market outcomes for urban couples

	First: Number of Children	First: No Boy	Total Hours	Total Income	Formal Hours	Formal Income
45<age<60	-0.288*	0.332***	-53.959**	1514.873	-31.141*	-1236.795
	(0.170)	(0.087)	(27.161)	(5004.408)	(16.353)	(5324.569)
Gender first born	0.176***	0.601***				
	(0.021)	(0.011)				
marriedbefore75	0.843***	-0.033				
	(0.043)	(0.022)				
marriedbefore85	0.331***	-0.089***				
	(0.040)	(0.021)				
marriedafter00	-0.153	0.030				
	(0.133)	(0.069)				
Ethnic Majority	-0.121***	0.050**				
	(0.040)	(0.020)				
numberchildren			-1.499	-2181.313	-24.461***	-5228.034***
			(8.751)	(1612.361)	(5.269)	(1715.513)
No Boy			-6.555	715.062	4.689	1098.204
			(6.359)	(1171.566)	(3.828)	(1246.518)
45<age<60*No Boy			16.680	6439.858**	19.702**	10213.288***
			(15.956)	(2939.839)	(9.607)	(3127.917)
age>60*No Boy			-3855.979	32944.589	-2251.397	-456757.575
			(2563.352)	(472294.012)	(1543.338)	(502509.333)
45<age<60*Nr. Child			26.045	-2350.530	10.374	-1139.379
			(16.186)	(2982.244)	(9.745)	(3173.035)
age>60*Nr. Child			651.652	8412.745	400.765	81959.502
			(430.102)	(79245.694)	(258.955)	(84315.490)
cohort dummies	Yes	Yes	Yes	Yes	Yes	Yes
year dummies	Yes	Yes	Yes	Yes	Yes	Yes
province dummies	Yes	Yes	Yes	Yes	Yes	Yes
obs.	5030	5030	5030	5030	5030	5030

Notes: Same as Table 4 except that the sample is now on couples in urban areas.

Details on the implementation

The unknown parameters are estimated by indirect inference. Many papers, either explicitly or implicitly, follow this approach. Examples are Gourinchas and Parker (2002), French (2005), Attanasio, Low, and Sánchez-Marcos (2008), Van der Klaauw and Wolpin (2008), and French and Jones (2011).

The idea is that we simulate, for given parameters θ, data from the model

$$y_{it} = m(x_{it}, \varepsilon_{it}; \theta).$$

That is, for a given θ, we can simulate choices y_{it}^s, $s = 1, \ldots, S$. Our data are the actual choices individuals i made in periods $t = 1, \ldots, T$, given the observed x_{it}. Denote the distribution of the simulated data in simulation round s with NT observations, by $\hat{F}^s_{y_{it}|x_{it};\theta}$. We denote this function with a hat because it is the "empirical" distribution that is based on a finite number of observations.

There is an auxiliary model, with an auxiliary parameter β. If y_{it} is a binary choice, then the auxiliary model could for example be a parametric binary choice model. Or any other model for which we can find a criterion function Q such that the so-called binding function that was introduced by Gourieroux and Montfort (1995),

$$b(\theta) = \arg\max_{\beta} Q\left(\beta, F_{y_{it}|x_{it};\theta}\right),$$

can be defined. Here, $F_{y_{it}|x_{it};\theta}$ is the true distribution of y_{it} given the exogenous state variables in x_{it}, but not the endogenously determined ones, for the parameter value θ. Q could be a likelihood function or the negative of the generalized method of moments (GMM) objective function.[11] Assume that Q has a unique maximizer for any given θ. In our case, the auxiliary models will correspond to those underlying the four last columns in Table 4. Since we can control all factors leading to endogeneity in our model, we will use simple linear regressions instead of the instrumental variables estimators.

Define β_0 to be the value of the auxiliary parameter once we evaluate the binding function at the true value of the structural parameters, θ_0,

$$\beta_0 = b(\theta_0).$$

Assume, for identification, that $b(\cdot)$ is invertible and denote its inverse by $b^{-1}(\cdot)$. Moreover, assume that the Jacobian matrix $\partial b(\theta)/\partial \theta'$ is of full column rank when evaluated at the true parameter value θ_0.

At this point, it is useful to make the following thought experiment. Suppose $b(\theta)$ was known. Then, we could use data to estimate β_0, and use the parameter estimates

[11] It would have to be the negative of the GMM objective function because usually we minimize the GMM objective function.

$$\hat{\beta} = \arg\max_{\beta} Q\left(\beta, \hat{F}_{y_{it}|x_{it}}\right)$$

to obtain

$$\hat{\theta} = b^{-1}(\hat{\beta}).$$

However, this is generally not feasible because the binding function depends on the unknown $F_{y_{it}|x_{it};\theta_0}$. The idea of indirect inference is to instead use the distribution of the simulated data, $\hat{F}^s_{y_{it}|x_{it};\theta}$, over multiple simulation rounds s, to estimate θ.

In round s, and for each value of the structural parameters we can simulate data and then estimate the auxiliary parameters as

$$\hat{\beta}^s(\theta) = \arg\max_{\beta} Q\left(\beta, \hat{F}^s_{y_{it}|x_{it};\theta}\right).$$

Here, we implicitly use the empirical distribution function of y_{it} given x_{it} for the simulated data set of size NT. Asymptotically, that is when the number of observations in the original data, or in one simulated data set, goes to infinity, we have that $\hat{F}^s_{y_{it}|x_{it};\theta}$ goes to $F_{y_{it}|x_{it};\theta_0}$ and hence $\hat{\beta}^s(\theta) = b(\theta)$. The idea of the indirect estimator of θ is to calibrate the value of θ such that

$$\frac{1}{S}\sum_{s=1}^{S} \hat{\beta}^s(\theta)$$

is close to $\hat{\beta}$, which is the estimated parameter of the auxiliary model for the actual data.

Formally, it is the solution of a minimum distance problem and akin to the GMM estimator,

$$\hat{\theta} = \arg\min_{\theta} \left[\hat{\beta} - \frac{1}{S}\sum_{s=1}^{S} \hat{\beta}^s(\theta)\right]' \hat{\Psi}^{-1} \left[\hat{\beta} - \frac{1}{S}\sum_{s=1}^{S} \hat{\beta}^s(\theta)\right], \tag{1}$$

where $\hat{\Psi}$ is a positive definite matrix converging to a positive definite matrix Ψ. Under the aforementioned assumptions it is a consistent estimator of θ.

This shows that the estimator is obtained by evaluating $\hat{\beta}^s(\theta)$ only at specific values appearing in the optimization. It is based on simulation and therefore the same draws of ε_{it} are used throughout. As for the elements of x_{it} that are exogenous to the model, one always uses the ones from the original data.

For fixed S, the estimator is asymptotically normally distributed with

$$\sqrt{NT}(\hat{\theta} - \theta_0) \xrightarrow{d} \mathcal{N}(0, \Lambda),$$

where

$$\Lambda = \left(1+\frac{1}{S}\right)\,(D'\Psi^{-1}D)^{-1}\,(D'\Psi^{-1}\Phi\Psi^{-1}D)\,(D'\Psi^{-1}D)^{-1} \tag{2}$$

with

$$D = \left.\frac{\partial b(\theta)}{\partial \theta'}\right|_{\theta=\theta_0}$$

and

$$\Phi = J_0^{-1}(I_0 - K_0)J_0^{-1}$$

for some I_0, J_0 and K_0. $I_0 - K_0$ can consistently be estimated by

$$N \cdot \frac{1}{S}\sum_{s=1}^{S}(W_s - \bar{W})(W_s - \bar{W})'$$

with

$$W_s = \left.\frac{\partial Q\left(\hat{\beta}, \hat{F}^s_{y_{it}|x_{it};\tilde{\theta}}\right)}{\partial \beta}\right|_{\beta=\hat{\beta}}$$
$$\bar{W} = \frac{1}{S}\sum_{s=1}^{S} W_s.$$

where $\tilde{\theta}$ is a consistent estimator for θ, for instance $\hat{\theta}$. J_0 can be consistently estimated by

$$\left.\frac{\partial^2 Q\left(\beta, \hat{F}_{y_{it}|x_{it}}\right)}{\partial \beta \partial \beta'}\right|_{\beta=\hat{\beta}},$$

the Hessian of the objective function of the auxiliary model evaluated at $\hat{\beta}$.

From (2), we can see that the optimal choice of the weighting function is

$$\left(\Psi^{-1}\right)^* = J_0(I_0 - K_0)^{-1}J_0,$$

assuming that $I_0 - K_0$ is invertible. Then,

$$\begin{aligned}\Lambda^* &= \left(1+\frac{1}{S}\right)\,\left(D'\left(\Psi^{-1}\right)^* D\right)^{-1}\,\left(D'\left(\Psi^{-1}\right)^*\Phi\left(\Psi^{-1}\right)^* D\right)\,\left(D'\left(\Psi^{-1}\right)^* D\right)^{-1} \\ &= \left(1+\frac{1}{S}\right)\,(D'\left(\Psi^{-1}\right)^* D)^{-1}.\end{aligned}$$