Essays in Household Saving Behavior and Effects of Growth

by

Noam Gruber

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2012

Doctoral Committee:

Professor Jonn P. Laitner, Chair
Professor Alan V. Deardorff
Associate Professor Dmitriy L. Stolyarov
Assistant Professor Jagadeesh Sivadasan

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER I

Keeping up with the Zhangs – household saving behavior as a function of relative wealth

"In any community where goods are held in severalty it is necessary, in order to his own peace of mind, that an individual should possess as large a portion of goods as others with whom he is accustomed to class himself; and it is extremely gratifying to possess something more than others."

– Thorstein Veblen, The Theory of the Leisure Class (1899)

1.1 Introduction

It is well established in economic literature, from Kuznets (1953) to Dynan et al. (2004), that the rich save a larger share of their income. In this study, I use Chinese household survey data to show that not only do the rich save more, but wealth, or richness, is not absolute, and is perceived relative to local average income levels. That is, households in high average income locations are likely to consume more of their disposable income, *ceteris paribus*. Using city identifiers, I estimate local average income, and construct a given household's relative income as the ratio of the household's income over the local average. I then demonstrate that the household's saving rate is more closely correlated to its relative income than to its absolute income. I am able to show this both for household reported income and, using education as an instrument, for the household permanent income / lifetime wealth.

To account for this relativistic household behavior, I construct a hybrid utility function that incorporates relative consumption, also referred to as "keeping up with the Joneses", and utility from wealth, also called the "capitalist spirit" after Weber (1905). Through analysis and simulations, I show that this hybrid utility generates higher rates of saving by the relatively rich. In addition, it is consistent with finding of high aggregate saving rates in rapidly growing economies.

The relative role of consumption, dating back to Veblen (1899) and the Duesenberry (1949) relative income hypothesis, has long been noted in the field of happiness economics.[1] Harbaugh (1996) and Carroll et al. (1997) suggest that relative consumption may explain the positive correlation between economic growth and saving (although their analysis is using an infinite horizon setup, which greatly weakens the effect). With economic growth, the surrounding average level of consumption is constantly rising, and the agent's optimal consumption path is exponentially increasing. This trend places the bulk of consumption later in life, when income is low (due to retirement), thus necessitating more saving earlier. However, a relative consumption model by itself cannot explain why the relatively rich save more, as it generates identical saving rates across different levels of income. Similarly, utility from wealth by itself does not account for the fact that it is the *relative* wealth of a household that determines its average propensity to consume or save, as neither heterogeneity in time discounting nor income mean reversion suffice to explain the magnitude of savings. Consequently, I argue that combining utility from relative wealth with relative consumption presents a more comprehensive explanation of saving behavior.

To my knowledge, this paper is the first to show increased saving rates with *relatively* higher income, and the first to combine relative consumption and utility from wealth. This 'hybrid' utility function closely matches the relation between relative income and the saving rate, as well as the stylized fact of a positive correlation between aggregate growth rates and saving to GDP ratios.

Section 1.2 presents the puzzles this paper addresses. Section 1.3 describes the survey data used and provides a definition of household saving. Section 1.4 documents the empirical relation found between relative income and the household saving rate. Section 1.5 analyses the relation between income and wealth, and presents the effect of estimated relative wealth on the household saving rate. Section 1.6 includes a discussion of various possible explanations to the aforementioned relation. Section 1.7 analyzes joining together utility from relative consumption and utility from relative wealth in order to replicate the observed saving behavior. Section 1.8 compares simulations of a model using such preferences to the data. Section 1.9 concludes.

[1] A seminal paper, presenting the 'Easterlin paradox', is Easterlin (1974). More recent papers discussing relative consumption are Luttmer (2005), Clark et al. (2008) and Oshio et al. (2010).

1.2 The Puzzles

1.2.1 The Relatively Rich Save Relatively More

That the rich (households in the upper income percentiles) save a larger share of their income is an accepted phenomenon in economic research. Kuznets (1953) shows it to be an extremely consistent in the US during the first half of the 20th century. Dynan et al. (2004) demonstrate it for recent US data. Chamon and Prasad (2010) show a similar pattern for Chinese households. Out of multiple explanations suggested, Carroll (2000), Dynan et al. (2004) and Francis (2009) find the most likely answer to be a form of utility from wealth, much like the treatment of wealth as a luxury good.

However, utility from wealth by itself fails to explain a crucial fact - that it is the relative, and not the absolute, wealth that determines the saving rate. As I show in Section 1.4, two households with the same income could have very different saving rates, depending on their relative income in their respective localities. I document this fact cross-sectionally using Chinese urban household and city-level data. Anecdotally, the relative nature of wealth and saving is also supported by the observation that while real income has seen a substantial increase in many countries over the second half of the 20th century, the saving rate in these countries has not increased.[2] Saving rates seem to be influenced by the rate of development, but not by its level.

1.2.2 Economic Growth and Household Saving Behavior

The phenomenon of high saving rates accompanying rapid growth is well documented in the economic literature. Modigliani (1970) notes this relationship in the context of post World-War-II Europe. More recently it has been observed in the rapidly growing Asian economies, most notably China. As Deaton and Paxson (2000) state, a one percent increase in per-capita growth is accompanied by a two percent increase in the saving rate.[3]

Despite empirical evidence of the positive relationship between economic growth and saving rates, standard infinite horizon representative agent and life cycle models predict the opposite. That is, such models suggest that higher expected growth leads to greater borrowing, or consuming at the expense of future wealth. In the

[2]Notice the similarity to the Easterlin paradox, in which richer households are happier, but happiness is not higher in richer countries.

[3]For empirical support for the relation between economic growth and the aggregate saving rate, see Carroll and Weil (1994) and Andersson (1999).

international context, these models suggest that a rapidly-developing country will run trade deficits, essentially borrowing from rich countries. However, the opposite is often observed in reality.[4]

In addition to its effect on aggregate and individual saving rates, rapid growth has two more closely related effects. First, Carroll and Summers (1991) find that households in rapidly growing countries experience a higher rate of increase in their consumption expenditure compared to their peers in lower growth countries. This phenomenon, observed in Lee et al. (2006), is inconsistent with the standard permanent income hypothesis (PIH) and life cycle models. The second phenomenon, observed in Carroll and Summers (1991), is that while different cohorts can have dramatically different levels of lifetime wealth, this does not cause a large difference in their consumption expenditure. For example, the lifetime (expected) wealth of a 60 year old is much lower than that of a 30 year old in a rapid growth environment, with this difference smaller in a slow growth location. Hence in a rapidly growing economy, we would expect a 60 year old to consume considerably less than a 30 year old. Interestingly, these differences do not seem to exist in data comparing US and Chinese expenditure data, as seen in Figure 1.1.

1.3 The Data: China Urban Household Survey, Years 1993-1997

My data, acquired from the Universities Service Centre for China Studies in the Chinese University of Hong-Kong, contains detailed income and expenditure cross-sectional information on roughly 31,000 Chinese households, repeated from 1993 to 1997 (about 6,200 households per year), spread over 57 cities, across ten of China's provinces.[5] The survey is performed by the Urban Socio-Economic Survey Organization of the Chinese State Statistical Bureau.

Household are chosen through a random process. Income and expenditure data are collected monthly. Surveyors are instructed to cross-check income and expenditure data using receipts and bank reports. Surveyed households are required by Chinese law to respond truthfully and accurately, and the privacy of their responses is protected by law. If a household is having difficulties filling out the monthly reports, due to a low educational level or other reasons, the surveyors are required to

[4] An example in point is Chinese savings flowing into the US.

[5] An overview of this dataset can be found online at http://www.usc.cuhk.edu.hk/DCS/DCS31-1-93-97.aspx.

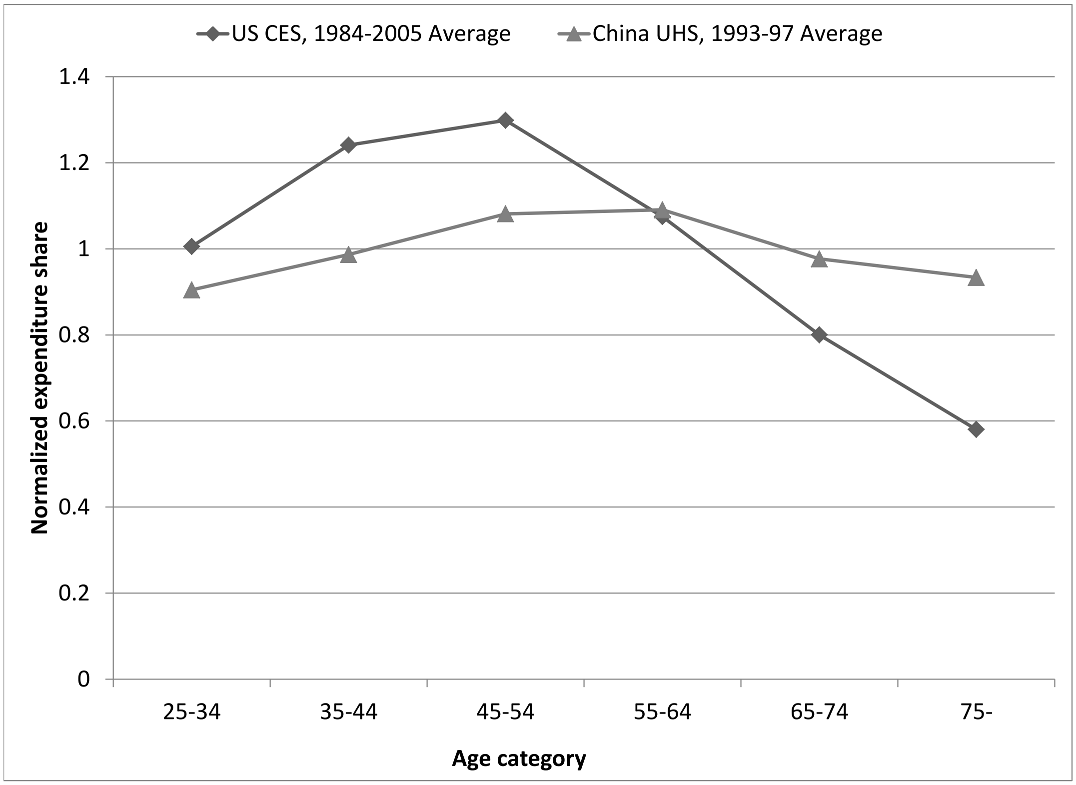

Figure 1.1: US and China Consumption-Age Cross-sections

visit multiple times per month in order to assure accurate reporting. Collected data is inspected at the city, the provincial and the country levels for misreporting and/or surveyor negligence.

On a side note, the household selection process should capture illegal rural residents in these urban locations, and the survey forms indeed allow for different residency permit (Hukou) types. Theoretically there is no reason for such households to decline participation in the survey, as they are protected by Chinese survey privacy laws. In practice, it is impossible to ascertain whether their participation rates are comparable to those of the legal residents.

The different cities provide significant heterogeneity in average income, both cross-sectionally and over time. Using this income heterogeneity, I demonstrate the effect of local average income on a given household's expenditure level.

Key to this analysis is the inclusion of city identifiers, which enables me to estimate average income by locality. While this feature is unobtainable from many US household level datasets, it is available in this Chinese data.[6]

[6]For example, the US Consumer Expenditure Survey contains, at best, state-level identifiers,

The availability of city identifiers and the significant average income heterogeneity between cities makes this data set an excellent testing ground for the effects of relative income. An added advantage to using Chinese data is the Chinese public record system (Dang'an tixi), which poses significant barriers to endogenous migration. The main disadvantages of this dataset are the lack of comprehensive household asset data, preventing both a direct measure of household wealth and the inclusion of unrealized capital gains, and the lack of longitudinal household data, preventing the estimation of household fixed effects.

As city-level price data is unavailable for China, I use provincial-level price indices to deflate some of the cost-of-living difference between the city-year pairs.

1.3.1 Defining Saving

To begin the analysis, the term 'saving rate' needs to be formally defined. This can be done in one of two ways. First, saving can be defined as unconsumed income, so that the saving rate is

$$s_i = 1 - \frac{CE_i}{Y_i}, \tag{1.1}$$

where Y and CE are disposable income and consumption expenditure, respectively, of household i. I use standard definitions of income and consumption: Disposable income includes work income (including bonuses and benefits), capital income and incoming transfer payments (such as pensions, subsidies and family assistance). Consumption expenditure includes food, clothing, furniture, appliances, decorations, utensils, home services (e.g., cleaning, cooking and child-sitting), medical expenses, communication, transportation, recreation, education and childcare, rent and other commercial goods and services. It does not include interest, loans paid, gifts, family assistance, house purchases or building expenses.

Second, saving can be defined as explicit expenditure on saving vehicles, so that the saving rate is

$$s_i = \frac{SV_i}{Y_i}. \tag{1.2}$$

where SV is the net household expenditure on saving vehicles. These include deposits and withdrawals to and from saving accounts and saving type insurance accounts, loans taken and made, bought and sold securities, credit or advance payment

while the American Community Survey (ACS) Public Use Microdata Sample (PUMS) includes geographical identifiers, but no expenditure data.

purchases, expenditure on purchasing or building a house, change in cash in hand, investment in household secondary business production and private transfers.[7] To reiterate, since I do not have data on asset stocks, unrealized capital gains are not included.

Fortunately, in this dataset the above two measures are virtually identical (they exhibit a correlation of over 0.99). This would seem to be a product of the access given to surveyors to household income, expenditure and financial reports, and of the use of the balance of income and expenditure for the supervision of surveyors' diligence.

I use the measure of saving based on expenditure on saving vehicles, as in Equation 1.2. In addition, I divide saving into three subcategories:

1. Financial saving, including net deposits to saving accounts and saving type insurance accounts, securities, loans, credit or advance payment purchases

2. Housing investment, including housing purchase and construction expenses as well as housing loans made and repaid

3. Private transfer payments, including financial support and gifts given out to family and friends[8]

Household heads ages censored to 25 to 70 to maintain good support. To avoid household composition endogeneity issues (e.g., the choice of adult children to stay or establish their own household), household income is defined as the sum of the incomes of the household head and spouse.

1.4 Income and the Saving Rate

As mentioned, it is well-documented that the wealthy save a larger share of their income. I will first analyze the effect of household income on the saving rate. Income is certainly not a perfect proxy for wealth, and the income-saving relation may be substantially different from the wealth-saving relation. This important issue is discussed in section 1.5.

Dynan et al. (2004) and Francis (2009) show the saving rate to be increasing in income, taken as a proxy for wealth. For China, this has been observed by Chamon and Prasad (2010), who show the difference in saving rates between income deciles in

[7] Private transfers constitute financial support and gifts given out. I follow Chamon and Prasad (2010) in categorizing these as saving.

[8] most notably to one's older relatives.

China to be increasing from 1995 to 2005. According to Chamon and Prasad (2010), in 2005, the average saving rate of the first income decile was around zero, while that of the tenth decile was about 30%.[9]

Figure 1.2 shows the saving rate to be increasing in income deciles computed for the 1993-1997 period. Decomposing savings into the three categories defined in Section 1.3.1, we can see the financial saving appears to generate most of the change in the overall saving rate. The rate of housing investment at the tenth decile is almost twice that of the first decile, but its effect is fairly small. The rate of private transfers seems flat across the income distribution.

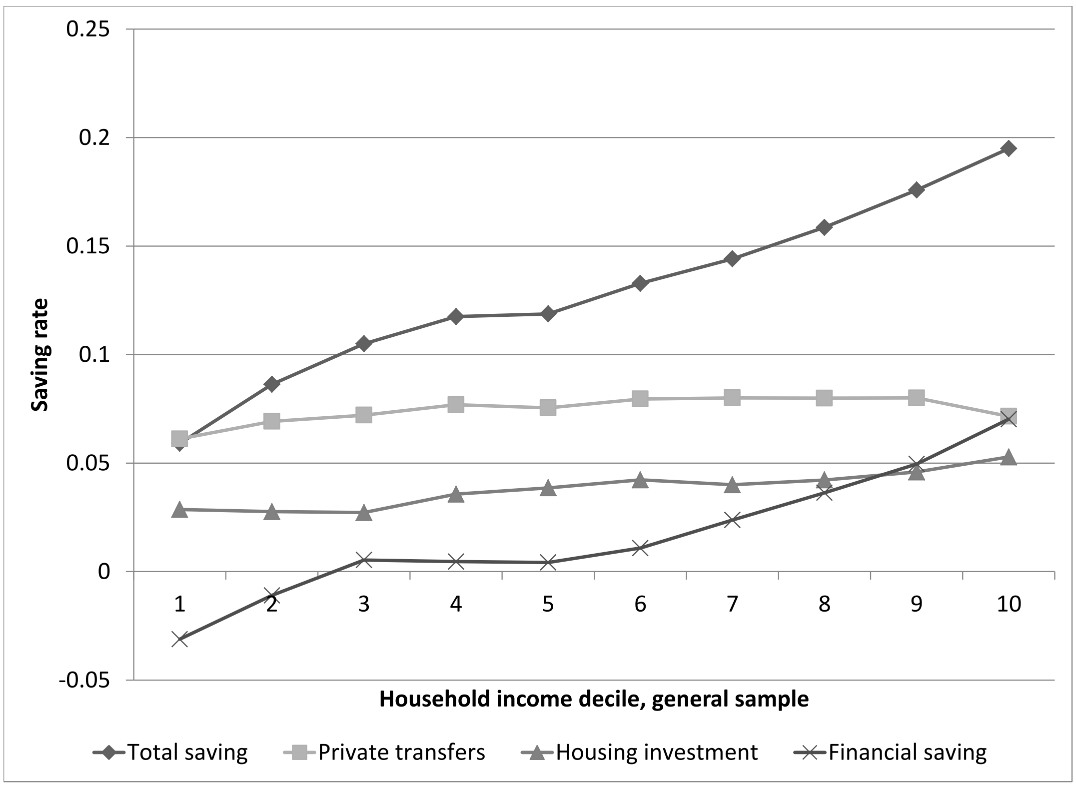

Figure 1.2: Saving rates by household income decile, general sample

Figure 1.3 shows this to be the case when computing income deciles for the single year of 1993.[10]

Next I examine the importance of relative income on saving choices. To graphically demonstrate the effect, I first divide the city-year pairs into quartiles based on the

[9]In view of this paper's results, a possible explanation for this increase in saving rate disparity is an increase in within-location income inequality over that period.

[10]1993 is chosen arbitrarily, the income-saving rate relation holds for all years in the sample.

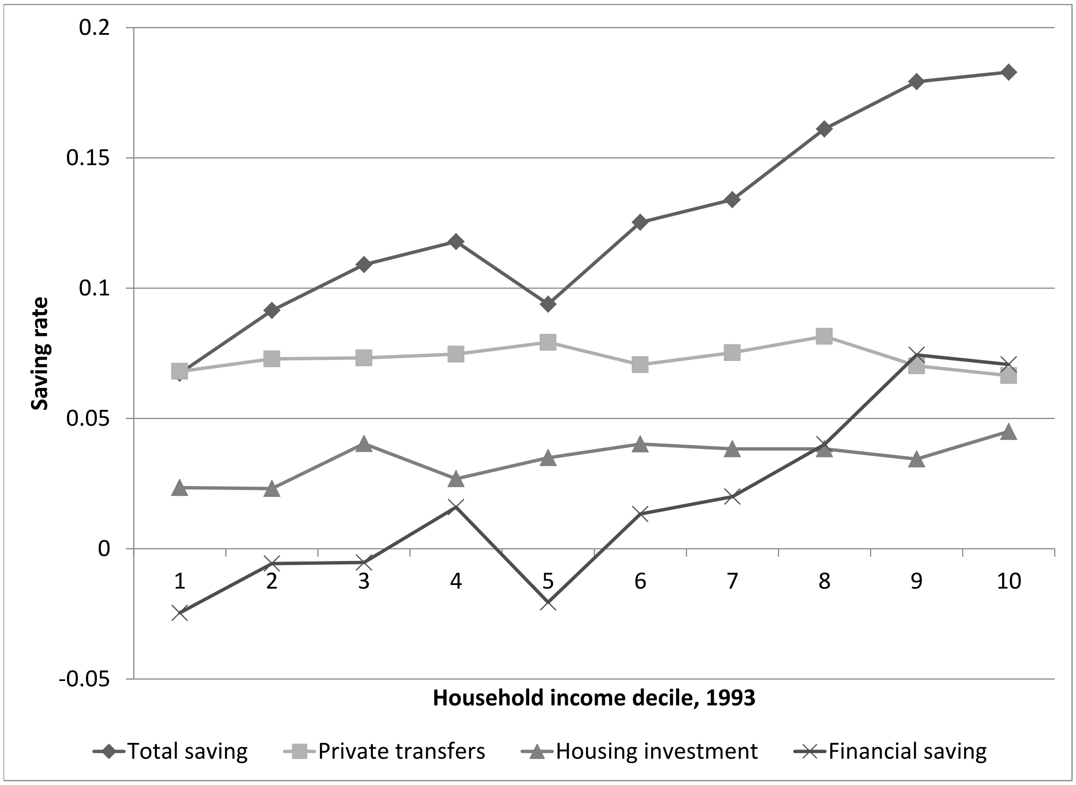

Figure 1.3: Saving rates by household income decile, 1993

average household income per city. Table 1.1 gives the mean income, saving rates and other descriptive statistics by income quartiles of the surveyed city-year pairs.

Average household income in the fourth quartile is more than double that of the first, but the saving rates are not very different. The ratio of private transfers and housing investment seem to be decreasing between income quartiles, while the financial saving rate is increasing. The average age of the household head is slightly higher in higher income city-year pairs, while the ratio of household heads employed by state owned enterprises (SOEs) is about the same. The share of academics (college educated or higher) and high-school educated household heads is increasing in city-year income, while the share of middle-school education is decreasing. Generally speaking, households in higher city-year income quartiles include more adult children and older relatives and less minor children.

A concern for such an analysis is that due to growth, later-year city-year pairs populate the higher quartiles, while early year pairs dominate the lower quartiles, so that the table largely reflects the time trends in the descriptive statistics, and not the differences between low and high income locations. To assuage this concern, Table

City-Year Income Quartile	**1**	**2**	**3**	**4**	**Total**
Household income (RMB per year)	6186.5	7623.6	8614.7	13100.4	8868.6
Saving rate	0.137	0.126	0.108	0.146	0.129
Private transfers (rate)	0.095	0.079	0.065	0.059	0.075
Housing investment (rate)	0.056	0.032	0.033	0.032	0.038
Financial saving (rate)	-0.015	0.015	0.010	0.055	0.016
Age of household head	44.4	44.9	45.6	46.3	45.3
State employee	0.693	0.673	0.650	0.667	0.671
Academic degree	0.056	0.070	0.086	0.098	0.077
Professional degree	0.123	0.130	0.144	0.141	0.135
Secondary school	0.153	0.144	0.126	0.116	0.135
High-school	0.219	0.211	0.240	0.246	0.229
Middle-school	0.333	0.346	0.324	0.304	0.327
# of adult children	0.429	0.411	0.407	0.470	0.429
# of minor children	0.779	0.706	0.675	0.692	0.713
# of older relatives	0.083	0.085	0.103	0.119	0.098

Table 1.1: Descriptive statistics (means) by city-year income quartiles, years 1993-1997

1.2 displays the same statistics for year 1993 only.

Overall Table 1.2 presents a very similar picture. Average household income in the fourth quartile is still almost double that of the first, but now the average saving rates of the first quartiles is the highest. The ratio of private transfers and housing investment still seem to be decreasing between income quartiles, while the financial saving rate is still increasing. The average age of the household head is now highest in the third income quartile, while the ratio of household heads employed by state owned enterprises (SOEs) is now decreasing in quartiles. The shares of academics and high school graduates show no definite trends. The number of adult children, minor children and older relatives also shows no trend.

To summarize Tables 1.1 and 1.2, households in the different city or city-year income quartiles have (by construction) very different incomes, on average, but are mostly comparable otherwise. The average saving rate in higher income cities or city-year pairs does not seem to be systematically higher.

Considering the Chinese high national saving to GDP ratio, the average household saving rate in Tables 1.1 and 1.2 may seem low. However, it is important to remember that since high saving rates are correlated with high income, the aggregate household saving rate is understated by averaging saving rates across households. In other words, the average of the saving rates is lower than the average saving rate. It is also

City Income Quartile	**1**	**2**	**3**	**4**	**Total**
Household income (RMB per year)	5841.9	7262.0	8162.0	11516.4	8043.3
Saving rate	0.152	0.112	0.123	0.119	0.126
Private transfers (rate)	0.108	0.071	0.058	0.053	0.073
Housing investment (rate)	0.043	0.035	0.037	0.020	0.034
Financial saving (rate)	0.000	0.005	0.028	0.045	0.018
Age of household head	43.6	43.9	47.5	45.5	45.0
State employee	0.717	0.698	0.665	0.617	0.677
Academic degree	0.056	0.077	0.121	0.083	0.083
Professional degree	0.108	0.112	0.141	0.097	0.115
Secondary school	0.147	0.129	0.124	0.130	0.133
High-school	0.221	0.212	0.193	0.257	0.220
Middle-school	0.355	0.366	0.315	0.316	0.341
# of adult children	0.477	0.380	0.447	0.479	0.442
# of minor children	0.830	0.738	0.649	0.763	0.746
# of older relatives	0.085	0.090	0.101	0.099	0.093

Table 1.2: Descriptive statistics (means) by city income quartiles, 1993

important to note that, as Deaton (2005) argues, household surveys often undersample high income households.[11] In addition, the aggregate saving rate in China has further increased over the years since the survey used in this work was taken.

Figure 1.4 shows graphically the importance of relative income, rather than absolute income, on the household saving behavior. First, I define the household's relative income as the ratio of its income to the average income in its location. Next I sort all households into relative income deciles. Grouping households by the city-year income quartiles detailed in Table 1.1, Figure 1.4 shows that households of very different absolute income exhibit, on average, similar saving saving behavior if their relative income is similar.

Once again, to avoid the suspicion that the similarity between the saving behavior of household in different city-year income quartiles is due to growth in income over time (in essence, each quartile could be the same group of cities in a different year), Figure 1.5 shows the same phenomenon for year 1993 alone.

Table 1.3 presents the results of regressing the household saving rate on different combinations of the log of household income, the log of average household income in the locality, the log of relative income (the difference between the two) and demographic covariates. All regressions are clustered by city.

[11] Chamon et al. (2010) discuss the inconsistency between Chinese household survey and national accounts data for saving rates over time.

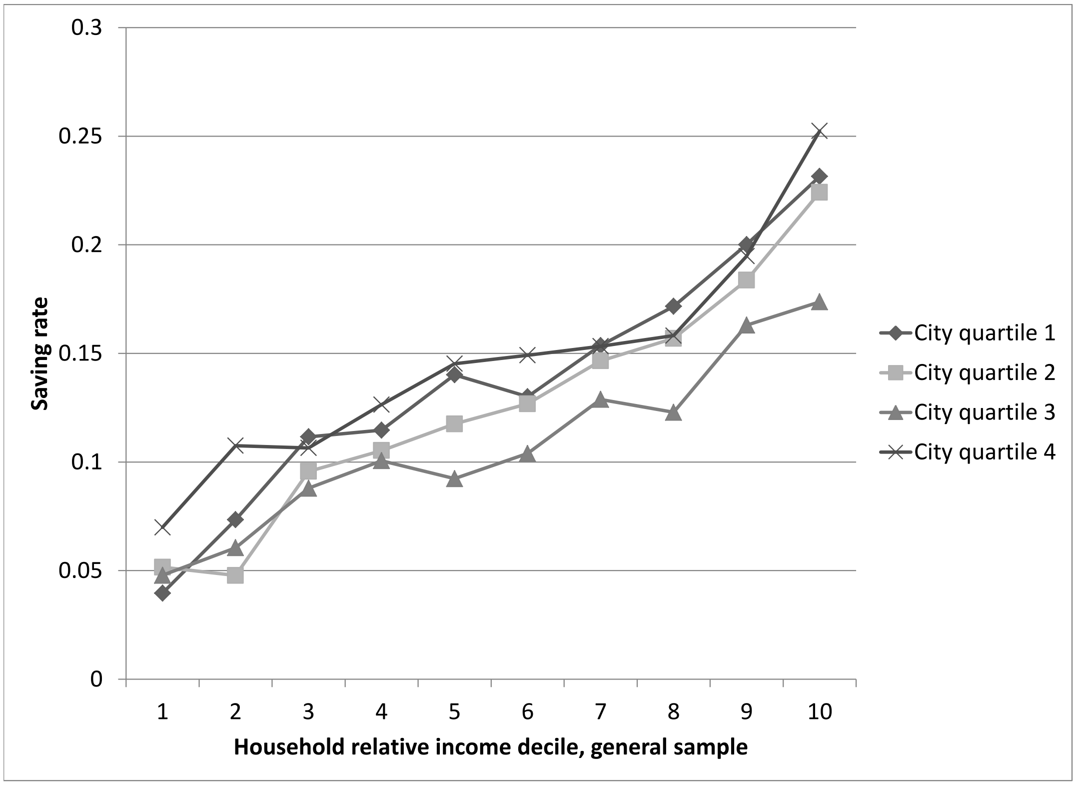

Figure 1.4: Saving rates by relative income deciles for city-year average income quartiles, general sample

As expected, column 1 in Table 1.3 shows that (absolute) household income has a positive effect on that household's saving rate. With a t-statistic over 6, this result is significant. Column 2 introduces the log of average income. With a t-statistic over 20, this result is also significant. Its inclusion causes the t-statistic of household income to increase from 6 to 10, and the R^2 to increase by about 50%, from 0.045 to 0.064. As expected, the sign of the coefficient on local average income is negative, meaning that higher local average income leads households to consume more and save less. Note that the size of the respective effects of income and average income is similar, so that the hypothesis that it is only the relative income that affects the household's saving rate cannot be rejected at a 5% statistical significance level.

	Absolute	Abs. and Local	Abs. and Relative	Relative	City FE
	(1)	(2)	(3)	(4)	(5)
log income	.096	.162	.014		
	(.015)***	(.008)***	(.013)		
log average city income		-.148			
		(.015)***			
log relative income			.148	.161	.163
			(.015)***	(.008)***	(.008)***
age of head	-.015	-.012	-.012	-.012	-.007
	(.003)***	(.003)***	(.003)***	(.003)***	(.003)***
age squared	.0005	.0003	.0003	.0003	.0002
	(.0001)***	(.0001)***	(.0001)***	(.0001)**	(.0001)
age cubed	-4.00e-06	-2.25e-06	-2.25e-06	-2.09e-06	-4.21e-07
	(1.39e-06)***	(1.44e-06)	(1.44e-06)	(1.45e-06)	(1.26e-06)
academic head	-.002	-.018	-.018	-.017	-.016
	(.010)	(.009)**	(.009)**	(.009)*	(.007)**
academic spouse	-.027	-.039	-.039	-.038	-.034
	(.011)**	(.012)***	(.012)***	(.012)***	(.010)***
professional school head	.0005	-.016	-.016	-.016	-.015
	(.012)	(.011)	(.011)	(.011)	(.009)*
professional school spouse	-.019	-.031	-.031	-.030	-.029
	(.008)**	(.008)***	(.008)***	(.008)***	(.008)***
secondary school head	.015	-.003	-.003	-.003	-.005
	(.008)*	(.007)	(.007)	(.007)	(.006)
secondary school spouse	-.017	-.030	-.030	-.029	-.031
	(.008)**	(.008)***	(.008)***	(.008)***	(.007)***
high-school head	.004	-.006	-.006	-.006	-.006
	(.008)	(.008)	(.008)	(.008)	(.007)
high-school spouse	-.010	-.014	-.014	-.013	-.015
	(.007)	(.007)**	(.007)**	(.007)*	(.006)**
middle-school head	.005	-.004	-.004	-.004	-.004
	(.007)	(.006)	(.006)	(.006)	(.005)
middle-school spouse	-.007	-.011	-.011	-.011	-.007
	(.007)	(.007)*	(.007)*	(.007)	(.006)
male	.007	.006	.006	.006	-.005
	(.008)	(.008)	(.008)	(.008)	(.004)
self-employed	.035	.031	.031	.031	.033
	(.015)**	(.015)**	(.015)**	(.015)**	(.014)**
adult children	.021	.024	.024	.024	.020
	(.004)***	(.003)***	(.003)***	(.003)***	(.002)***
minor children	-.023	-.024	-.024	-.024	-.028
	(.005)***	(.004)***	(.004)***	(.004)***	(.004)***
home owner	.018	.019	.019	.019	.0008
	(.008)**	(.008)**	(.008)**	(.008)**	(.005)
Const.	-.592	.146	.146	.265	.240
	(.131)***	(.122)	(.122)	(.026)***	(.020)***
Obs.	28847	28847	28847	28847	28847
R^2	.045	.064	.064	.063	.097
F statistic	22.893	50.46	50.46	52.627	54.297

Table 1.3: Saving rate, affected by absolute, local and relative income, in logs (clustered by city). Year dummies used but not displayed

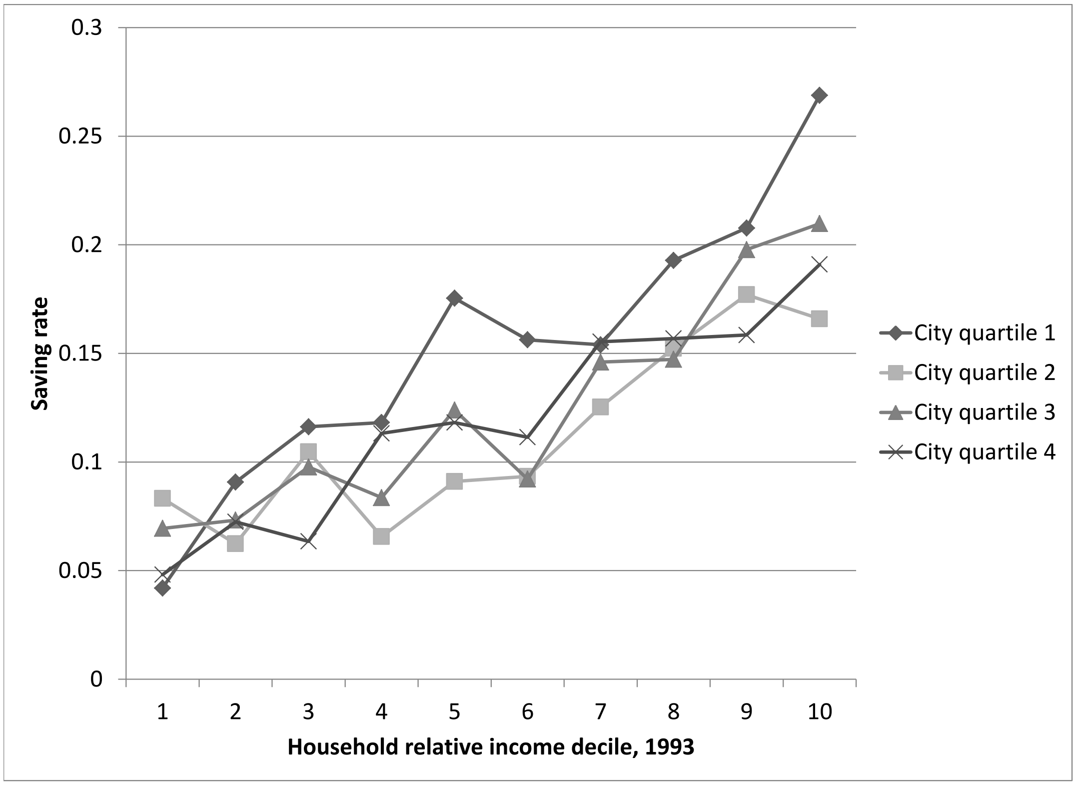

Figure 1.5: Saving rates by relative income deciles for city average income quartiles, 1993

Column 3 includes household income and relative income. In this regression, absolute income is no longer statistically significant. Column 4 drops absolute income. The R^2 is virtually unchanged. The coefficients of the rest of the variables are hardly affected. In column 5 I include city fixed effects, to absorb any city specific conditions, such as different price levels. The R^2 increases to 0.097. City fixed effects are not included in the regressions presented in columns 1-3 because they would absorb the effect of average city income and make it impossible to discern between relative and absolute income.

To summarize, columns 1-5 indicate that, *ceteris paribus*, it is the relative income of a household compared to the local average that determines its saving rate, and that absolute income has no additional explanatory power.

Furthermore, secondary, professional and academic education generally seem to lower the saving rate. The gender of the household head has no significant effect. Households in which the head is self-employed have a higher saving rate, possibly due to higher income volatility. Having adult children in the household seems to increase the saving rate, while minors seems to lower it. The home ownership coefficient is positive, but is not statistically significant in column 5.

Figure 1.6 displays the income shares of the three components of saving, namely financial saving, private transfers and housing investment, across relative income deciles. While financial savings still play a large role, here, unlike in Figure 1.2, housing investment and private transfers are also substantial contributors to the increase in the overall saving rate.

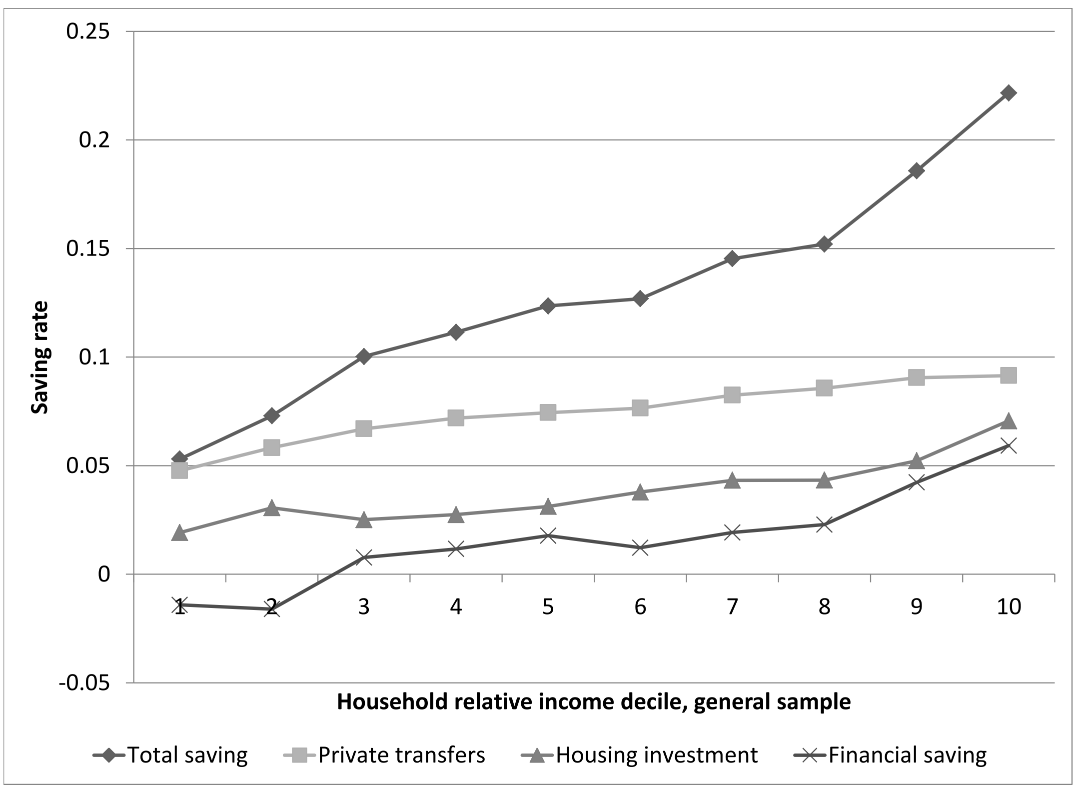

Figure 1.6: Saving rates by household relative income decile and by saving type, general sample

This again is mirrored using 1993 data alone, in Figure 1.7.

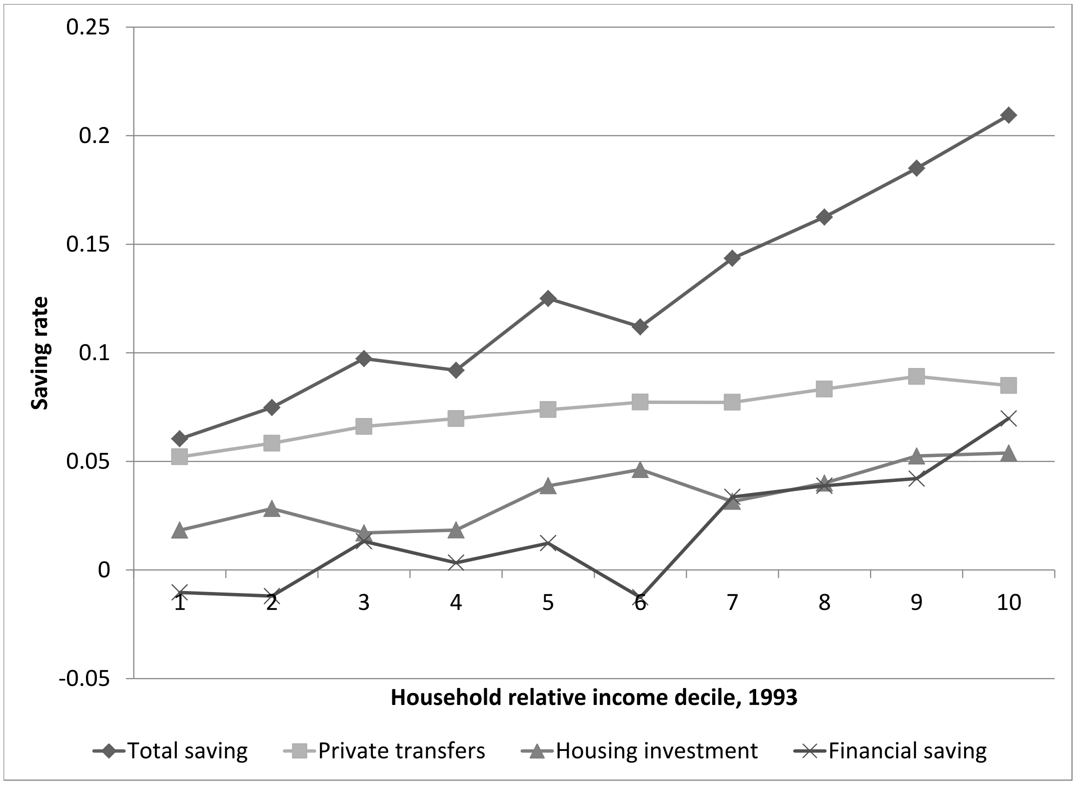

Figure 1.7: Saving rates by household relative income decile and by saving type, 1993

Table 1.4 presents estimation results for the three saving subcategories. The effects of individual household income and local average income do not seem to cancel out for these components of saving. The null hypothesis that the coefficients of log income and log average income sum up to zero is easily rejected.

For private transfers and housing investment the negative effect of average city income is much larger than the positive effect of individual household income, while for financial saving the individual effect is much larger and the average city income effect is not even statistically significant.

These results suggest that private transfers and housing investment, saving ve-

	Private Transfers	Housing Investment	Financial Saving
	(1)	(2)	(3)
log income	.028	.036	.097
	(.003)***	(.006)***	(.011)***
log average city income	-.067	-.055	-.025
	(.012)***	(.009)***	(.022)
log relative income			
age of head	-.006	-.0006	-.006
	(.001)***	(.002)	(.004)
age squared	.0003	.00002	.00004
	(.00005)***	(.00009)	(.0002)
age cubed	-3.41e-06	-1.93e-07	1.38e-06
	(5.44e-07)***	(1.13e-06)	(1.94e-06)
academic head	.007	.011	-.035
	(.003)**	(.008)	(.013)***
academic spouse	-.001	-.002	-.036
	(.005)	(.009)	(.017)**
professional school head	.014	.007	-.038
	(.004)***	(.007)	(.013)***
professional school spouse	.006	-.009	-.027
	(.005)	(.007)	(.011)**
secondary school head	.014	.004	-.020
	(.003)***	(.007)	(.009)**
secondary school spouse	.008	-.003	-.034
	(.003)***	(.006)	(.012)***
high-school head	.010	-.005	-.011
	(.003)***	(.007)	(.010)
high-school spouse	.003	-.007	-.010
	(.003)	(.005)	(.010)
middle-school head	.006	-.009	-.002
	(.002)***	(.006)	(.008)
middle-school spouse	.004	-.003	-.011
	(.003)	(.005)	(.009)
male	-.002	.005	.002
	(.002)	(.003)	(.009)
self-employed	-.009	.003	.037
	(.004)**	(.025)	(.030)
adult children	-.005	-.002	.031
	(.001)***	(.002)	(.004)***
minor children	-.013	-.002	-.009
	(.002)***	(.003)	(.006)
home owner	.017	.018	-.017
	(.004)***	(.006)***	(.009)*
Const.	.457	.208	-.524
	(.108)***	(.071)***	(.176)***
Obs.	28847	28847	28847
R^2	.079	.009	.02
F statistic	40.678	5.622	23.298

Table 1.4: Private transfers, housing investment and financial saving rates, affected by absolute and local average income (in logs, clustered by city). Year dummies used but not displayed

hicles popular in low-income locations, are replaced by financial saving vehicles in high-income locations. One possible explanation is that households in more modern cities have better access to the financial system. An additional factor could be the relative weakening of unofficial family and social networks in these modern cities.

Figures 1.8, 1.9 and 1.10 reflect the estimation results in Table 1.4 in a graphical way. All three forms of saving increase with household relative wealth, but private transfers are lower for high city-year average income quartiles, whereas financial saving rates are higher. This also explains why in Figures 1.2 and 1.3, using absolute income deciles, private transfer and housing investment rates appear much flatter, and financial saving rates much steeper, than in Figures 1.6 and 1.7, using relative income deciles.

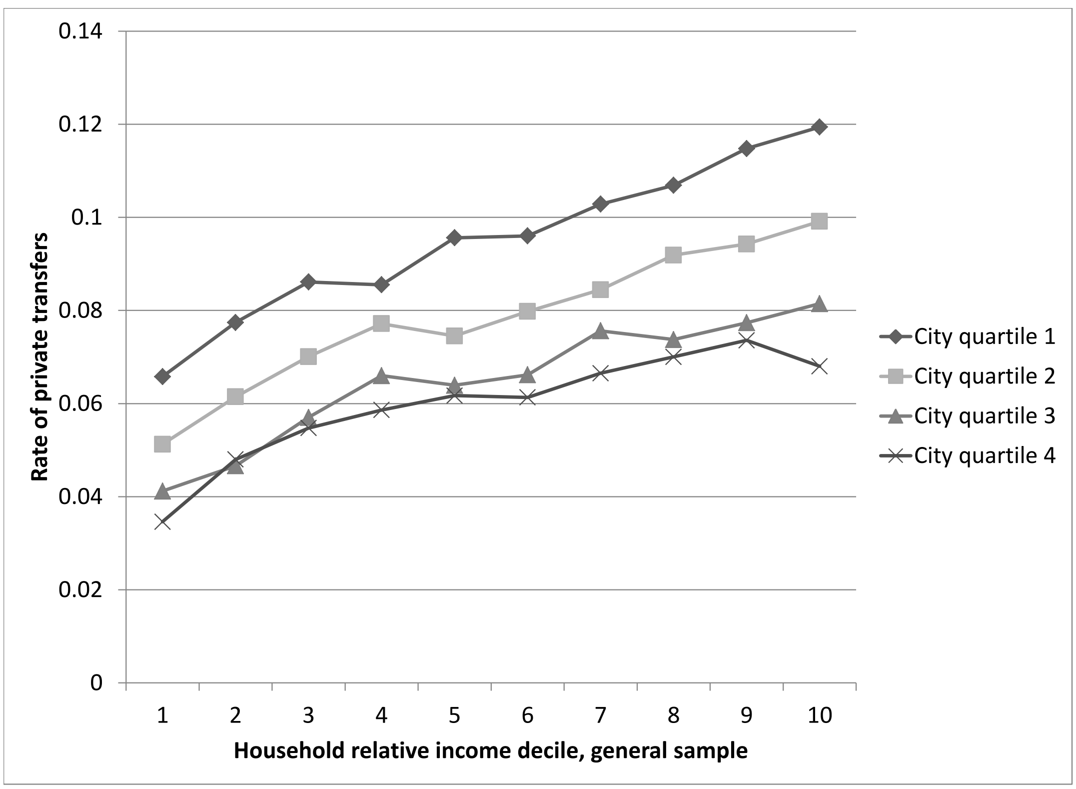

Figure 1.8: Private transfer rates by household relative income decile and city-year income quartiles, general sample

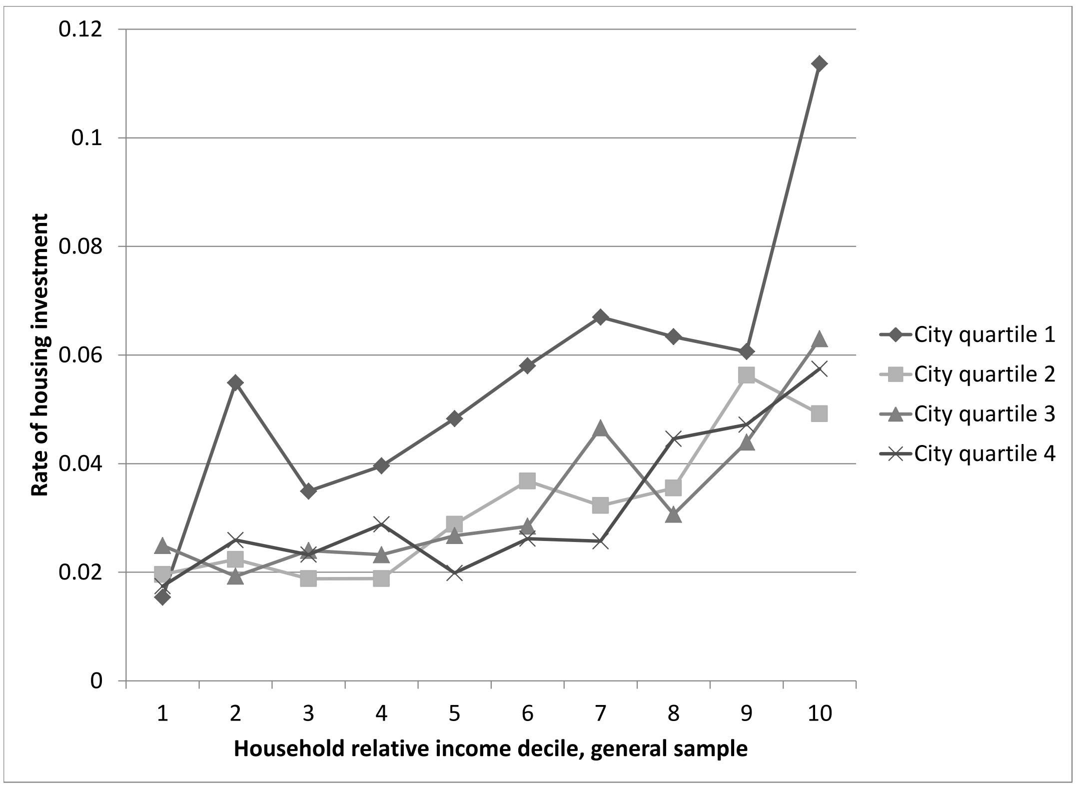

Figure 1.9: Rates of housing investment by household relative income decile and city-year income quartiles, general sample

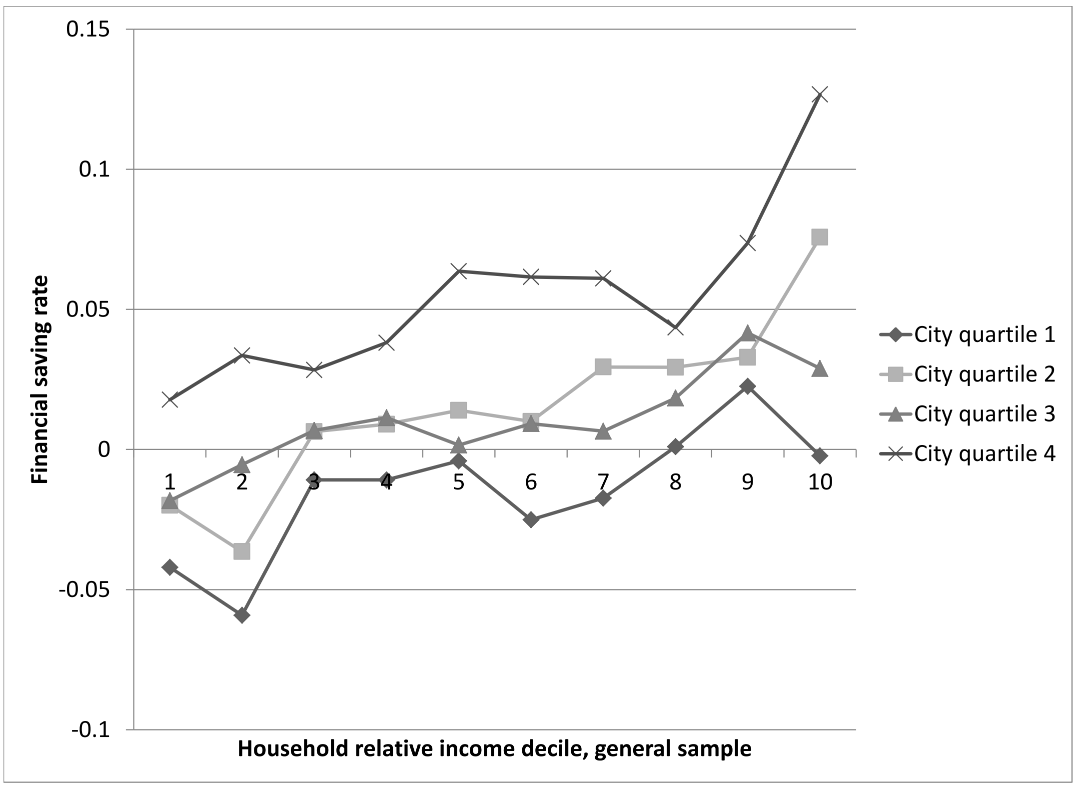

Figure 1.10: Financial saving rates by household relative income decile and city-year income quartiles, general sample

1.5 From Income to a Measure of Wealth

Breaking down surveyed income to its systematic components, it becomes evident that certain issues need to be addressed before income can be used as a proxy for wealth, especially in the context of measuring the effect of wealth on the saving rate.

In decomposing income I follow notation from King and Dicks-Mireaux (1982). First, permanent income is modeled as follows:

$$y_i^p = \boldsymbol{X_i'B} + s_i - c(age_i), \tag{1.3}$$

where y_i^p is permanent income, in logs, $\boldsymbol{X_i}$ is a vector of observables for household i, $\boldsymbol{B}$ is the vector of parameters, s_i is a mean zero household fixed effect and $c(age_i)$ is a cohort effect. Then y_{it}, yearly earnings in logs, is a function of permanent income, age and a deviation term,

$$y_{it} = y_i^p + h(age_{it} - age_0) + u_{it}, \tag{1.4}$$

where $h(\quad)$ represents the age-earning profile, and u_{it} is the deviation term. Let us break u_{it} into two components: transitory income, y_{it}^{tr}, and measurement error, me_{it}

$$u_{it} = y_{it}^{tr} + me_{it}. \tag{1.5}$$

I assume y_{it}^{tr} and me_{it} are not correlated. Combining equations 1.3, 1.4 and 1.5, we get

$$y_{it} = \boldsymbol{X_i'B} + g(age_{it}) + y_{it}^{tr} + s_i + me_{it}, \tag{1.6}$$

where $g(age_{it}) = h(age_{it} - age_0) - c(age_{it})$.

The decomposition of income to its structural components allows for a clear analysis of some of the problems in estimating the effect of income on the saving rate. First, a measurement error, in this case the misreporting of income, would create a spurious correlation between reported income and the saving rate. In the case of regressing reported income on the saving rate as defined in Equation 1.1, the spurious relation would be positive; if using Equation 1.2, the measure actually used here, the spurious relation would be negative. Second, if, as Friedman (1957) argues, consumption levels are largely determined by permanent income, then using an income measure that aggregates both permanent and transitory elements would create a correlation between income and the saving rate even if permanent income has no effect on it. Third, $c(age_i)$, the cohort effect, is hard to estimate because of the relatively small timespan of the data (5 years). Finally, the household fixed effect cannot be directly estimated, as household data is not longitudinal.

Since the data does not include repeated observation for the same household, the household fixed effect cannot be directly estimated. While King and Dicks-Mireaux (1982) suggest a procedure to artificially generate this effect, this would involve using parts of the deviation term, reintroducing a correlation between the error term and the saving rate. While having a good estimate for the mean zero household fixed effect would have improved the accuracy of the permanent income measure, omitting it should not bias the results.

The deviation term, mostly due to the transitory income part, is likely to be correlated with the saving rate. I take an instrumental variable approach to tackle this problem. Similarly to Dynan et al. (2004) I use educational dummies as the excluded variable. An age polynomial is used to absorb the effects of both the cohort and the age-earning profile.

As Table 1.5 shows, even when using an instrumental variable approach in order to approximate permanent income, the saving rate is increasing in individual (permanent) household income and decreasing in local average income. As with reported

	OLS	IV
	(1)	(2)
log income	.162	.072
	(.008)***	(.020)***
log average city income	-.148	-.064
	(.015)***	(.023)***
log relative income		
age of head	-.012	-.011
	(.003)***	(.003)***
age squared	.0003	.0004
	(.0001)***	(.0001)***
age cubed	-2.25e-06	-3.15e-06
	(1.44e-06)	(1.42e-06)**
academic head	-.018	
	(.009)**	
academic spouse	-.039	
	(.012)***	
professional school head	-.016	
	(.011)	
professional school spouse	-.031	
	(.008)***	
secondary school head	-.003	
	(.007)	
secondary school spouse	-.030	
	(.008)***	
high-school head	-.006	
	(.008)	
high-school spouse	-.014	
	(.007)**	
middle-school head	-.004	
	(.006)	
middle-school spouse	-.011	
	(.007)*	
male	.006	.006
	(.008)	(.007)
self-employed	.031	.033
	(.015)**	(.016)**
adult children	.024	.021
	(.003)***	(.003)***
minor children	-.024	-.023
	(.004)***	(.004)***
home owner	.019	.020
	(.008)**	(.008)***
Const.	.146	.156
	(.122)	(.119)
Obs.	28847	28847
R^2	.064	.049

Table 1.5: Comparing OLS and IV estimation - Saving rate, affected by absolute and local average income, in logs (clustered by city). Year dummies used but not displayed. Educational dummies used as instruments in IV.

income, both effects are of similar size and opposite sign. If permanent income can be considered a good measure of lifetime wealth, then this result suggests that the saving rate is a function of relative, rather than absolute, wealth.

1.6 The Hypothesis: Relative Consumption and Utility from Wealth

As noted in the introduction, the effect of comparative consumption on utility and happiness has long been a subject of research, with strong evidence that surrounding consumption levels affect individual utility. The present paper suggests that the elements of preference manifested in such behavior, namely relative consumption and utility from wealth, together explain differential saving rates. While each element already exists in the literature, this work is first to integrate the two into one model.

1.6.1 Reasons for Comparative Consumption Behavior

One intuitive explanation for this relative behavior is the notion of status seeking, as portrayed by Veblen (1899). According to this explanation, individuals and households, when their survival is assured, seek to maximize utility from the social status that consumption and / or asset holdings bestow upon them.[12] In this context, underconsumption (compared to the local norm) may indicate low earning ability and social inadequacy. Conversely, overconsumption of luxury goods may be pointless in a 'poor' environment,[13] where it could be viewed as vain and reckless, and might make one a target for crime. In an example of how consumption choices may influence social status, Cole et al. (1992) and Wei and Zhang (2009) argue that wealth, or the amount of one's savings, determines status and is influential in important non-market

[12] Assuming wealth is somewhat observable, this explanation bears some resemblance to the modeling of conspicuous consumption in Charles et al. (2009).

[13] If no one else has a car, owning a basic one will give you just about the same 'rich' status as buying a Rolls-Royce.

matters such as the choice of spouse. While such behavior may be suboptimal from an individual 'rationality' perspective , Weber (1905) would argue that it is optimal at the society level.

In this paper I focus on local average consumption as the household's base of self-comparison. In another study, Charles et al. (2009) shows that ethnic groups are also a powerful reference for determining one's comparative status.

I would like to note that there may be reasons other than status behind the importance of relative consumption.

First, some studies suggest that strong local social networks can affect household saving decisions, in fear of friends and relatives asking for financial help.[14]

Second, a local market is most likely geared toward common consumption. In order to consume rare goods, one may have to pay high premiums. In addition, information about such goods is scarce – local knowledge cannot be relied upon to decide on specific goods. Thus, individuals may retain local consumption patterns regardless of income.

Third, if consumption choices are at least partly formed by imitation, it is reasonable to assume similar consumption profiles by locality.[15] In other words, due to the practice of local imitation, wealthier households may not know how to adjust the composition of their consumption in order to extract high utility from its increase, and may thus reach something resembling satiation.

Fourth, average local income can conceivably be highly correlated with the level of local infrastructure, which in turn may a consumption facilitator. For example, the lack of paved roads would inhibit the purchase of vehicles, the lack of Internet access would diminish the attractiveness of computers and the lack of restaurants

[14]An example of a paper portraying such behavior is Goldberg (2010).

[15]It would seem that social networks propagate consumption patterns. For example, Christakis and Fowler (2007) find that obesity spreads over social networks, Christakis and Fowler (2008) find that smoking (and quitting smoking) is a network-based trait, and Rosenquist et al. (2010) show the same for alcohol consumption.

would limit dining out.

The above elements can all be part of the explanation for the role of relative income in determining household saving behavior. I choose to be agnostic about the weights of the roles played by different elements, and instead focus on what they all have in common: the importance of relative income in determining saving behavior.

1.6.2 Household heterogeneity

Reversing the causality, it could be argued that it is not that the rich save more, but rather that those who save more, become rich. In essence, this explanation suggests that there is something different about certain households which causes them to both earn and save more. Thus, extra saving may enhance their earning ability (through financing private business, obtaining higher education etc.). An oft-mentioned form of such heterogeneity is the level of time discounting, or patience. That is, patient households save more early in their lifecycle. Additionally, patience and ability may be correlated, causing patient households to have a higher income on average. However, both Carroll (2000) and Dynan et al. (2004) find that patience heterogeneity cannot fully account for the increased saving rates of the rich. In addition, such household would show lower saving rates later in the lifecycle, which is not the case, as evident in Figure 1.11.

1.6.3 Price levels

Given the existence of non-tradables, it stands to reason that price levels in high income locations are relatively high. One could argue that, due to different price levels, a household with a given income in a high income location actually has less purchasing power than a household of comparable income in a low income location. Thus it is not being poorer, relative to the local mean, that makes the former household save less, but rather actually being poorer, in terms of purchasing power. As

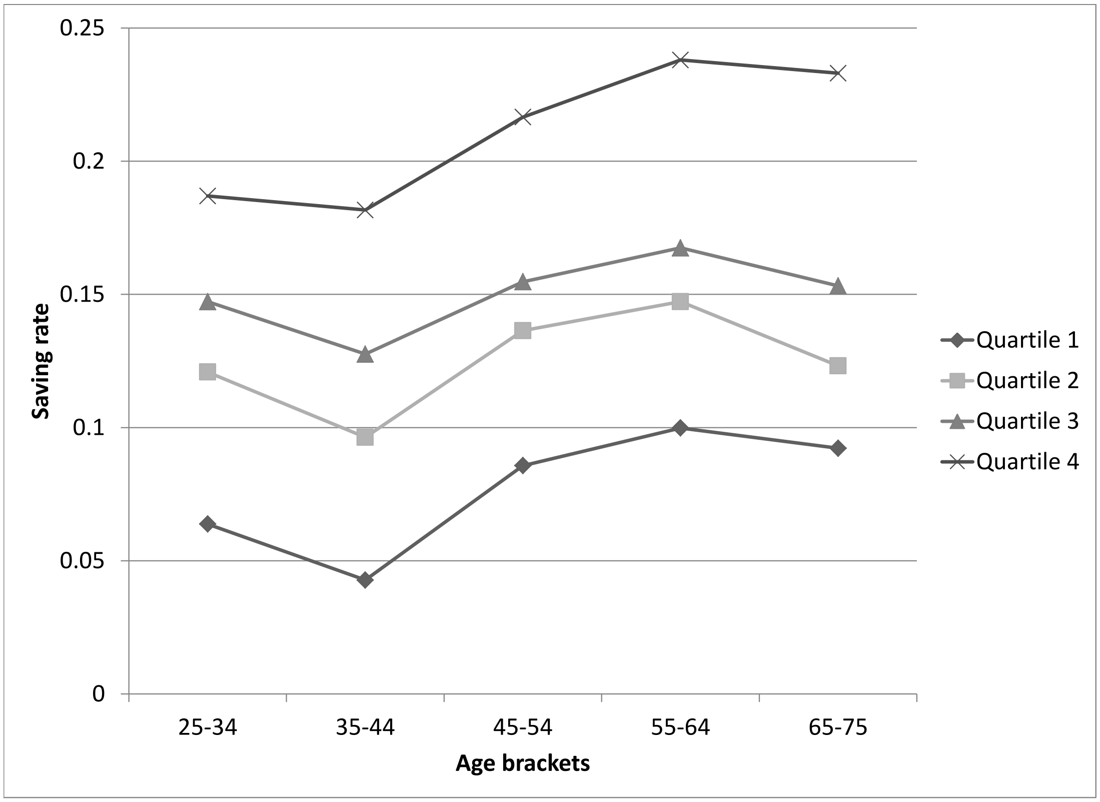

Figure 1.11: Saving rates over the lifecycle by household relative income quartiles (within age brackets)

Table 1.3 shows, the effects of household income and local average income on the saving rate are similar in size. This is also evident in Figure 1.4.

For the price level explanation to hold, prices would need to increase one-to-one with average local income. As long as prices do not increase one-to-one with average local income, the price level difference between localities does not weaken the explanatory power of relative income. To show this, let us first assume that P_i, the price level experienced by household i, is linearly correlated with local average income:

$$P_i = \phi \bar{y}_i, \tag{1.7}$$

where $\phi \geq 0$.

Then, subtracting the price levels (in log terms) from both household income and average local income, cancels out, as $\beta_1 + \beta_2 = 0$:

$$s_i = \beta_0 + \beta_1(y_i - P_i) + \beta_2(\bar{y}_i - P_i) + \beta_3 z_{1,i} + \epsilon_i = \beta_0 + \beta_1 y_i + \beta_2 \bar{y}_i + (\beta_2 - \beta_1) P_i + \beta_3 z_{1,i} + \epsilon_i. \tag{1.8}$$

On the other hand, if $\phi \approx 1$, the price level can replace local average income entirely:

$$s_i = \beta_0 + \beta_1(y_i - P_i) + \beta_2 z_{1,i} + \epsilon_i = \beta_0 + \beta_1 y_i - \beta_1 P_i + \beta_3 z_{1,i} + \epsilon_i \approx \beta_0 + \beta_1 y_i - \beta_1 \bar{y}_i + \beta_3 z_{1,i} + \epsilon_i. \tag{1.9}$$

If ϕ nears unity, it would indicate that given prices, the average purchasing power is equal across all cities. Nevertheless, the data and some back of the envelope calculations indicate that $\phi < 1$.

To illustrate, I identify the price level P_i is a function of the prices and shares of tradable and non-tradable goods. I then assume p_t, the price of tradables (uniform throughout the economy), $\bar{P}$, the national price level and $\bar{y}$, the national average income, are all normalized to unity. Next I assume that $p_{nt,i}$, the local price of non-tradables (mostly services and real-estate), is one-to-one with average income:[16]

$$p_{nt,i} = \bar{y}_i, \tag{1.10}$$

[16]This is a worst case scenario, as it is hard to justify prices going beyond one-to-one.

where $\bar{y}_i$ is relative to the national average $\bar{y}$.

It follows that:

$$P_i = (1 - \tau_i) + \tau_i \bar{y}_i, \tag{1.11}$$

where τ_i is the share of non-tradables at locality i.

Evidently, the local price level P_i can only be argued to near a one-to-one relation with average local income ($\phi \approx 1$) if τ_i is close to unity. However, using expenditure shares from the Chinese urban household survey as presented in Appendix 1.8, it is clear that the shares of food and clothing are consistently above 60% in all city groups. Therefore, τ, the share of services and other non-tradables, cannot approach unity. The consistency of expenditure shares also suggests that there is no major difference in relative prices between poorer and richer cities.

Consequently, it is unreasonable to assume that prices vary by the same scale as income, and thus a price-level explanation of the findings does not suffice. In addition, the price level explanation is insufficient to explain the fact that real income levels in developed countries have increased historically without a corresponding increase in saving rates.

To further illustrate this point, I perform the following exercise: I estimate the effect of absolute and local average income, both in OLS and IV form, for three categories of tenure choice – public housing (provided by the household's work unit), renting a house (a very uncommon choice in the sample) and home ownership. Since housing is a non-tradable good which can require a major expenditure during the household life cycle, we would expect that if the prices of non-tradables are driving the relation between relative wealth and the saving rate, this relation would be different for households with a different tenure choice. Table 1.6 presents the results.

	Public (OLS)	Rent (OLS)	Owner (OLS)	Public (IV)	Rent (IV)	Owner (IV)
	(1)	(2)	(3)	(4)	(5)	(6)
log income	.167	.176	.153	.068	.075	.071
	(.010)***	(.044)***	(.010)***	(.023)***	(.091)	(.035)**
log average city income	-.145	-.195	-.143	-.056	-.084	-.064
	(.019)***	(.047)***	(.023)***	(.025)**	(.084)	(.041)
age of head	-.009	-.009	-.005	-.011	-.042	-.010
	(.001)***	(.005)*	(.002)**	(.004)***	(.016)***	(.005)**
age squared	.0002	.0002	.0001	.0003	.002	.0004
	(.00002)***	(.0001)*	(.00004)**	(.0001)**	(.0007)**	(.0002)**
age cubed				-2.33e-06	-.00002	-4.34e-06
				(1.74e-06)	(9.36e-06)**	(2.47e-06)*
academic head	-.020	.0006	-.009			
	(.013)	(.084)	(.019)			
academic spouse	-.037	-.159	-.042			
	(.015)**	(.063)**	(.024)*			
professional school head	-.016	.001	-.022			
	(.015)	(.053)	(.014)			
professional school spouse	-.033	-.029	-.036			
	(.009)***	(.046)	(.017)**			
secondary school head	-.008	.075	.006			
	(.010)	(.054)	(.013)			
secondary school spouse	-.036	-.058	-.023			
	(.009)***	(.050)	(.014)			
high-school head	-.011	.011	.001			
	(.011)	(.049)	(.010)			
high-school spouse	-.012	-.018	-.024			
	(.009)	(.035)	(.012)**			
middle-school head	-.009	-.0006	-.003			
	(.009)	(.048)	(.010)			
middle-school spouse	-.009	-.033	-.011			
	(.009)	(.034)	(.011)			
male	.012	.032	-.010	.012	.029	-.008
	(.009)	(.031)	(.008)	(.009)	(.030)	(.008)
self-employed	.025	.039	.029	.029	.044	.030
	(.020)	(.043)	(.024)	(.021)	(.054)	(.025)
adult children	.027	.009	.025	.023	-.017	.021
	(.003)***	(.016)	(.006)***	(.003)***	(.020)	(.006)***
minor children	-.031	-.001	-.018	-.028	.008	-.016
	(.005)***	(.028)	(.006)***	(.005)***	(.030)	(.006)***
Const.	.066	.380	.129	.125	.458	.159
	(.172)	(.293)	(.168)	(.164)	(.248)*	(.171)
Obs.	18895	403	7833	18895	403	7833
R^2	.065	.136	.061	.047	.108	.047

Table 1.6: Saving rate by house ownership status (clustered by city). Year dummies used but not displayed

Households of different tenure choice exhibit a virtually identical relation between relative income (or wealth, as proxied for in the IV estimation). The low significance levels for the renter category (both with the OLS and IV estimations) are to be expected, given that this is a clustered regression with less than 10 observations per cluster, on average. In all regressions, the effect of local average income is of similar size and opposite sign to the effect of individual household income. It is still relative income that matters, regardless of tenure choice.

1.7 The Model: Relative Consumption and Utility from Wealth in a Lifecycle Setting

Any model attempting to replicate the empirical observation of the effect of relative income on the saving rate should ideally meet both of the following conditions:

Condition 1 (Saving rate neutral in equal relative income) *If* $\{s_t\}_{t=1}^{T}$ *is optimal given* $\{y_t\}_{t=1}^{T}$ *and* $\{\bar{y}_t\}_{t=1}^{T}$*, then* $\forall x > 0$, $\{s_t\}_{t=1}^{T}$ *is optimal for* $\{x * y_t\}_{t=1}^{T}$ *and* $\{x * \bar{y}_t\}_{t=1}^{T}$.

Condition 2 (Saving rate increasing in relative income) $\forall x > 1$, *if* $\{s_t\}_{t=1}^{T}$ *is optimal given* $\{y_t\}_{t=1}^{T}$ *and* $\{\bar{y}_t\}_{t=1}^{T}$*, and* $\{s_t^*\}_{t=1}^{T}$ *is optimal given* $\{x * y_t\}_{t=1}^{T}$ *and* $\{\bar{y}_t\}_{t=1}^{T}$*, then* $s_t^* > s_t \forall t$*,*

where s_t and y_t are the saving rate and income of the household in period t respectively, and $\bar{y}_t$ is the average income in the locality of the household in period t.

To create a hybrid utility function which meets the above conditions, I use elements from two types of models present in the literature: the spirit of capitalism models and the relative consumption models. Separately, each of these model types is unable to meet conditions 1 and 2. However, I show that by integrating the two model types, both conditions can be met.

Standard spirit of capitalism models, such as the ones featured in Carroll (2000), Dynan et al. (2004) and Francis (2009), replicate the increased saving rate of the rich by introducing assets into an intra-period utility function. Let

$$u(c_t, a_t) = \frac{c_t^{1-\theta}}{1-\theta} + \lambda \frac{(a_t + \chi)^{1-\alpha}}{1-\alpha}, \tag{1.12}$$

where c and a denote consumption and asset holding, respectively; θ and α are the consumption and asset coefficients of relative risk aversion, $\theta > \alpha$; λ is the weight on utility from assets, $\lambda \geq 0$; and $\chi \geq 0$ is the Stone-Geary luxury factor for asset utility.

Since this functional form treats assets like a luxury good, the share of assets increases with lifetime wealth. Since χ is (weakly) positive, low income households will hold assets only for lifecycle saving purposes (preferring utility from consumption over utility from assets). By contrast, high income households will hold assets for utility (overcoming χ and preferring to hold assets over the forgone consumption), with richer households expending a larger share of their wealth on assets, due to the $\theta > \alpha$ condition.

Note that the spirit of capitalism model fulfills Condition 2, since the saving rate increases in income, but fails to meet Condition 1, as average local wealth is not included in the model. Equation 1.12 implies that as average lifetime wealth increases, so should the aggregate saving rate.

Can a utility function using relative consumption account for households with relatively higher income saving more? Consider the following function, which uses the Galí (1994) formulation:

$$u(c_{t,i}, \bar{c}_{t,i}) = \frac{{c_{t,i}}^{1-\theta} \bar{c}_{t,i}^{\theta}}{1-\theta}, \tag{1.13}$$

where $c_{t,i}$ and $\bar{c}_{t,i}$ are individual and average household consumption in locality i at time t, respectively.

As Harbaugh (1996) and Carroll et al. (1997) discuss, a relative consumption element may generate a causal relation between economic growth and the saving rate. However, for a given time profile of average consumption, the inclusion of an average local consumption term amounts to multiplying a standard Constant Relative Risk Aversion utility function by a set of constants. As such, it does not cause high income households to save differently. Thus, it meets Condition 1 (as would any standard CRRA utility function), but fails to meet Condition 2.

Overall, neither element, by itself, is able to account for the observed phenomenon of relative income influencing a household's saving rate. However, the following intra-period hybrid utility function, is able to replicate this fact:

$$u(c_{t,i}, \bar{c}_{t,i}, a_{t,i}, \bar{a}_{t,i}) = \frac{c_{t,i}^{1-\theta} \bar{c}_{t,i}^{\theta}}{1-\theta} + \lambda \frac{\left(a_{t,i} + \chi \bar{a}_{t,i}\right)^{1-\alpha} \bar{a}_{t,i}^{\alpha}}{1-\alpha}, \tag{1.14}$$

where $\bar{a}_{t,i}$ is the average asset holding in locality i at time t.

In this model, relative assets function as a luxury good, whose share increases with relative wealth.[17] Note that, as in equation 1.13, this utility function is homogeneous of degree 1. Growth, also as in equation 1.13, increases both $\bar{c}$ and $\bar{a}$, thus causing an increasing consumption path for the household.

I now use equation 1.14 within a lifecycle perfect foresight framework:

$$U = \sum_{t=1}^{T} \beta^{t-1} u(c_{t,i}, \bar{c}_{t,i}, a_{t,i}, \bar{a}_{t,i}). \tag{1.15}$$

[17] The relative terms of consumption and assets can be interpreted as status indicators, as in Bakshi and Chen (1996).

Assuming $1+r=\frac{1}{\beta}$ for simplicity,[18] optimization yields the following Euler equations:

$$\left(\frac{c_{t,i}}{\bar{c}_{t,i}}\right)^{-\theta} = \left(\frac{c_{t+1,i}}{\bar{c}_{t+1,i}}\right)^{-\theta} + \lambda\left(\frac{a_{t,i}}{\bar{a}_{t,i}} + \chi\right)^{-\alpha} \quad \forall t < T, \tag{1.16}$$

$$\left(\frac{c_{T,i}}{\bar{c}_{T,i}}\right)^{-\theta} = \lambda\left(\frac{a_{T,i}}{\bar{a}_{T,i}} + \chi\right)^{-\alpha}. \tag{1.17}$$

Note that these Euler equations depend solely on consumption and asset ratios, and are neutral in levels.

As standard, the within period budget constraint is:

$$y_t + (1+r)a_{t-1} = c_t + a_t. \tag{1.18}$$

Finally, the household must have non-negative assets at the close of life:

$$a_T \geq 0. \tag{1.19}$$

Because households draw utility from their asset holdings, final period assets may be greater than zero. Since my model does not include utility from bequests, I assume all remaining assets are taken by the government, so that the income of existing households remains unaffected.

Other elements abstracted from the model include income uncertainty, timing-of-death uncertainty, private and public transfer payments (such as a pay-as-you-go pension), lifecycle weighted utility and household structure dynamics related to relative and absolute income levels (e.g., choice of number of children).

[18] The assumption that r is constant, consistent with a small open economy, is needed in order to make the proofs in this paper tractable. I further assume that $1+r=\frac{1}{\beta}$ for convenience, though this later assumption is not necessary for the proofs.

As the propositions below state, the hybrid utility function creates saving rate income level neutrality, so that only relative income affects the household saving behavior, and a household saving rate increasing with relative income. For the respective proofs, see Appendix 1.B.

Proposition 1 (Saving rate neutral in equal relative wealth) *If* $\{s_t\}_{t=1}^T$ *is optimal given* $\{y_t\}_{t=1}^T$ *and* $\{\bar{y}_t\}_{t=1}^T$*, then* $\forall x > 0$ $\{s_t\}_{t=1}^T$ *is optimal for* $\{x * y_t\}_{t=1}^T$ *and* $\{x * \bar{y}_t\}_{t=1}^T$.

Proposition 2 (Asset share for all periods increasing in relative wealth) $\forall x > 1$, *if* $\{a_t\}_{t=1}^T$ *is optimal given* $\{y_t\}_{t=1}^T$ *and* $\{\bar{y}_t\}_{t=1}^T$ *and* $\{a_t^*\}_{t=1}^T$ *is optimal given* $\{x * y_t\}_{t=1}^T$ *and* $\{\bar{y}_t\}_{t=1}^T$*, then* $\frac{a_t^*}{x*y_t} > \frac{a_t}{y_t} \forall t$.

Note that Proposition 1 is identical to Condition 1, whereas Proposition 2 is necessary, but not sufficient, for Condition 2. Analytically, I am able to prove only that the ratio of assets to income increases with relative income. While this does not prove that the saving rate increases with relative income, my proof implies that the ratio of final period assets, which is the discounted value of the stream of lifetime savings, to lifetime wealth is increasing in relative wealth.[19]

1.8 Calibration and Simulation

To test the ability of my model to replicate the empirical relation between relative wealth and the household's saving rate, I conduct a simulation. In this simulation, households live for 55 periods (roughly ages 25 to 80). The timing of household death and future income flows are both known in advance. I assume a constant growth rate throughout the household's lifetime.[20]

[19]In addition, my numerical results show a strong and consistent increase of the saving rate in relative wealth, as shown in Figures 1.18 and 1.16.

[20]This is a simplification, as high-growth countries are likely to catch up with the development frontier, slowing down as they near it.

Parameter	Value
T (length of household lifecycle)	55
θ (consumption utility coefficient)	3
α (assets utility coefficient)	1
λ (assets utility weight)	0.05
χ (assets utility Stone-Geary parameter)	0
β (time preference)	0.95
r (international real interest rate)	2%

Table 1.7: Calibration of Model Parameters

I use the income-age profile in my data, as seen in Figure 1.12, for the simulation. Of special interest is late-age income, typically comprised of pensions and other transfer payments. The relation between working-age and retirement-age income is of paramount importance in determining saving behavior in a lifecycle model. The survey data indicate a fairly flat post-retirement income, which, in a high growth environment, would suggest that transfer payments increase with economic growth.[21]

Another important issue in designing my simulation is the potential survey selection bias: low income elderly people are more likely to join their children's households, leaving the sample. I choose to use the household income-age profile in the data for calibration. An alternative would be to use an individual income-age profile, thereby circumventing the endogeneity of household formation and dissolution. However, the option of joining a child's household may be of great value to a parent, equivalent to substantial transfer payments. An individual income profile would not capture this value.

While my profile choice is explained above, the specific profile does not drive my results. I experiment with several profiles (e.g., a model with no pension income, a flat cross-section income profile up to retirement at household age 40 and a profile based on US data), and obtain similar qualitative results.

[21]While pensions are usually based on pre-retirement wage rates, and as such should not generally be affected by post-retirement growth, this does makes sense for private (within-family) transfers. If the income of working family members increases due to growth, they should be able to increase assistance for retired family members.

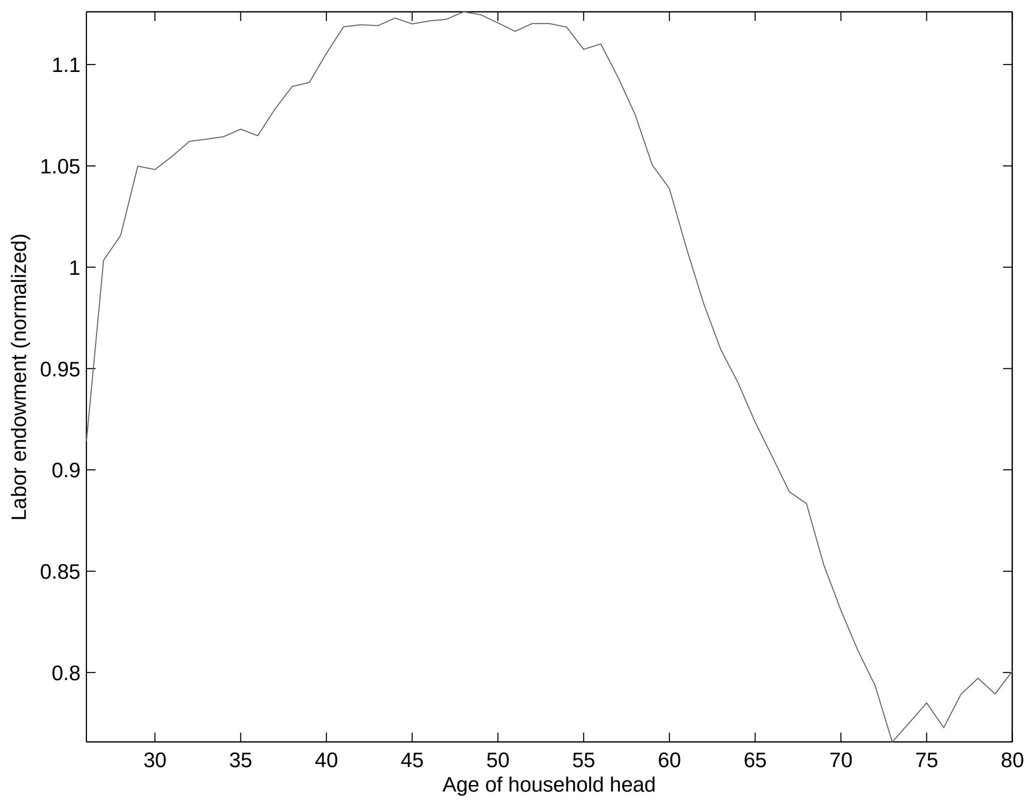

Figure 1.12: Labor endowment by age

To simulate income inequality, I draw income quantiles from a log-normal income distribution, with a mean and variance matching the those in the data. To reduce the effect of lifecycle income variation, I use only households with a head between 35 and 55 years old for calculating the income distribution's mean and variance. Income profiles are symmetric between quantiles, as illustrated in Figure 1.13.

I simulate a range of growth rates. Growth affects household income directly, as seen in Figure 1.14. This is the context in which late-age income matters most, since rapid growth has a very strong effect on it.

For θ, the consumption risk-aversion parameter, I use a value of 3. I set α, the asset semi risk-aversion parameter, to one third of θ, a value of 1, in order to replicate

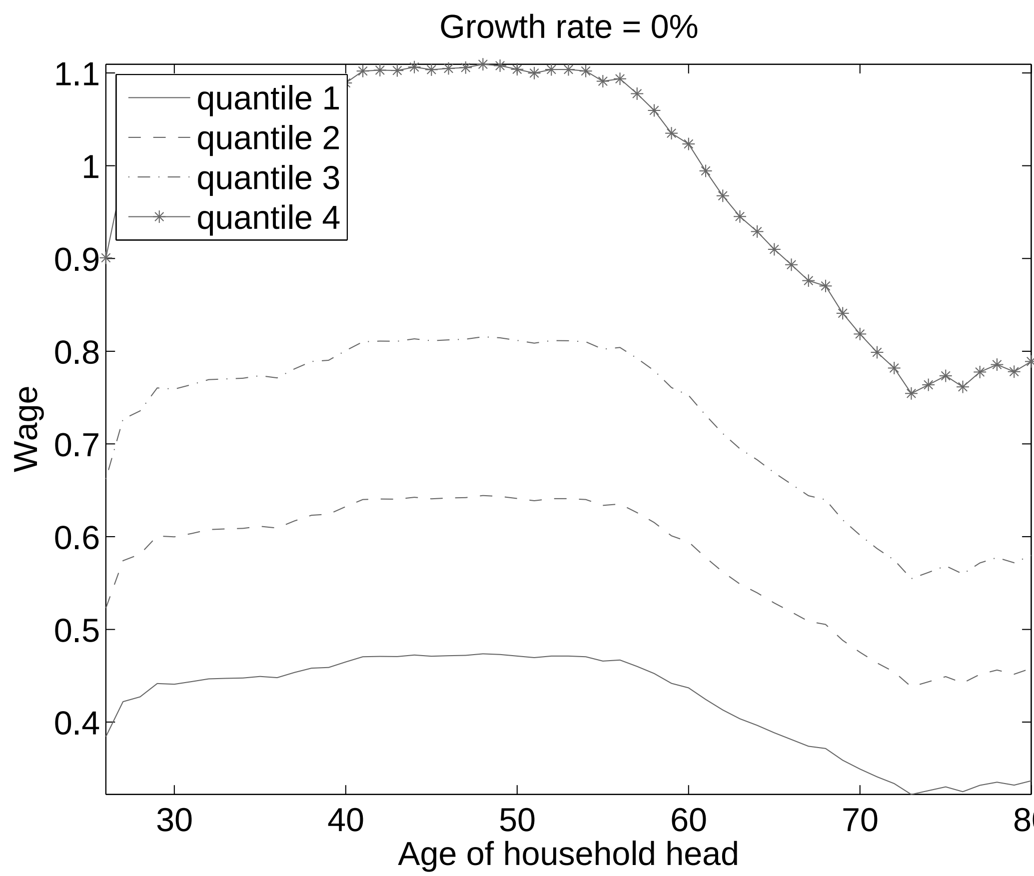

Figure 1.13: Income-age profile by quantile

the effect of greater asset holdings by the relatively rich as a share of lifetime wealth.

The remaining parameters, λ and χ somewhat offset each other. While λ encourages households to prefer asset holding to consumption and has a stronger effect on the rich, χ discourages households from holding assets, and has a stronger effect on the poor. However, since these effects are disproportional, they cannot be reduced to a single variable. While a positive λ is imperative for creating utility from assets, adding a positive χ stops lower income households from preferring asset holding to consumption, creating a saving-rate kink as described in Carroll (2000).

Thus, to simplify the calibration, I set χ to zero. While this means that all households hold assets for utility, implicitly entailing a non-negativity asset constraint

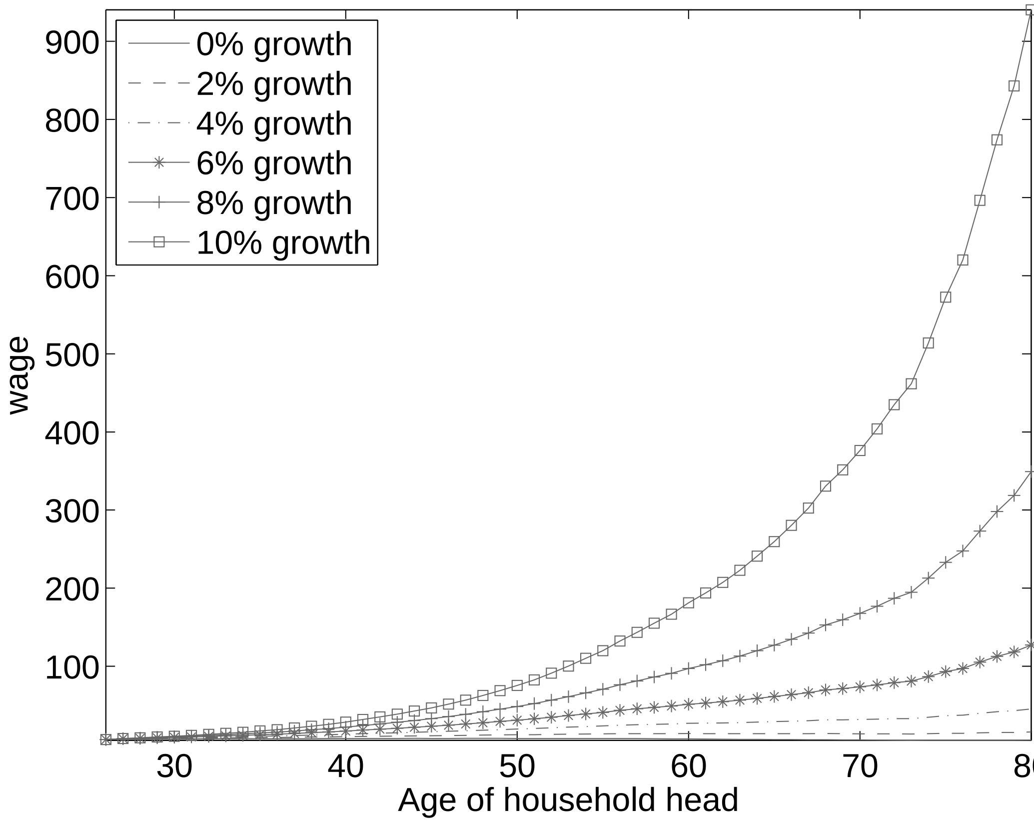

Figure 1.14: Average income-age profile by growth rate

and a positive end-of-life asset holding, the poor will hold very little of their lifetime wealth in assets while the rich will hold a much larger share. This is similar to findings from the US Survey of Consumer Finances data, as shown in Carroll (2000).

I set λ to a value of 0.05, to match the saving rate by income quantile profile from Figure 1.4.

1.8.1 Comparing Simulation and Data

I find that the simulation is able to replicate the main stylized facts from the data. The 4% growth line is quite similar to that of the Chinese data presented in Figure 1.4.

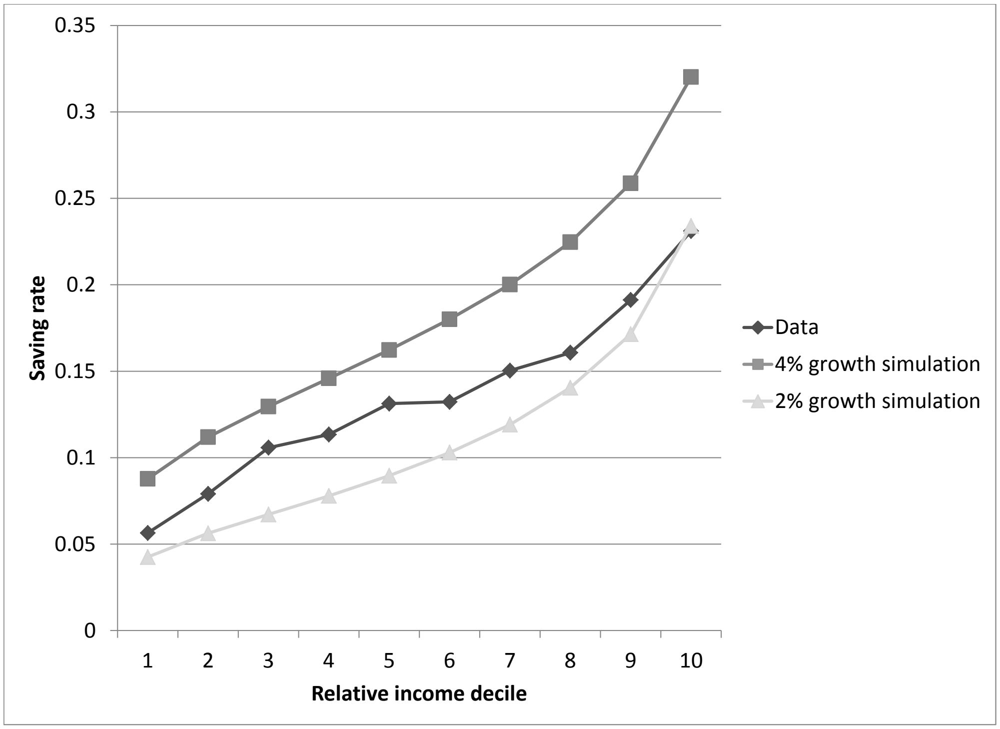

Figure 1.15: Saving rate by local income quantile

As in the data shown in Figure 1.11, the simulated saving rate differentials between income quantiles are consistent throughout the lifecycle. Figure 1.16 displays the saving rates across ages for different quartiles.

Simulation results are also consistent with the main stylized facts regarding the effects of growth on saving / consumption behavior. As Figure 1.17 shows, the consumption-age cross-section profile remains similar even across remarkably different growth rates, despite the inter-cohort lifetime wealth disparity created by growth.[22]

Finally, the simulation reinforces the point that the hybrid model is able to generate a positive relation between output growth and the aggregate saving rate. As Figure 1.15 shows, the saving rate increases for both income quantile and economic

[22]This conforms to the data, as discussed in Section 1.2.2.

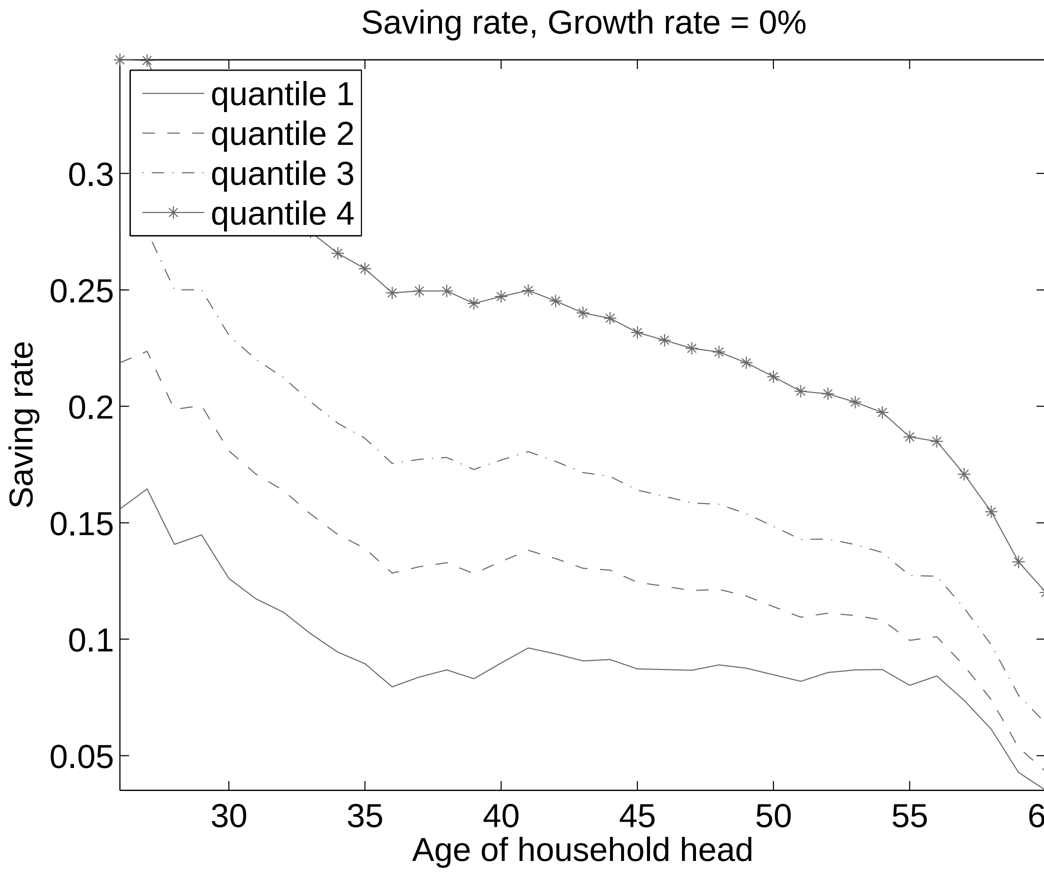

Figure 1.16: Saving rates over the lifecycle by quantiles

growth rate. Note that, for the zero growth line, the lower quantile households save almost nothing.

Figure 1.19 presents the simulated aggregate saving rate for a range of economic growth rates, compared with a corresponding international cross-section. The countries used for the international comparison are listed in Appendix 1.C. Growth rates are averaged over ten years, as they are much more volatile than the saving rate.

The simulated values of the aggregate saving rate by growth present a trend similar to the data, but are generally lower. This can be explained, at least partly, by the role of the other determinants of the aggregate saving rate, i.e., retained earnings by firms, and government savings.

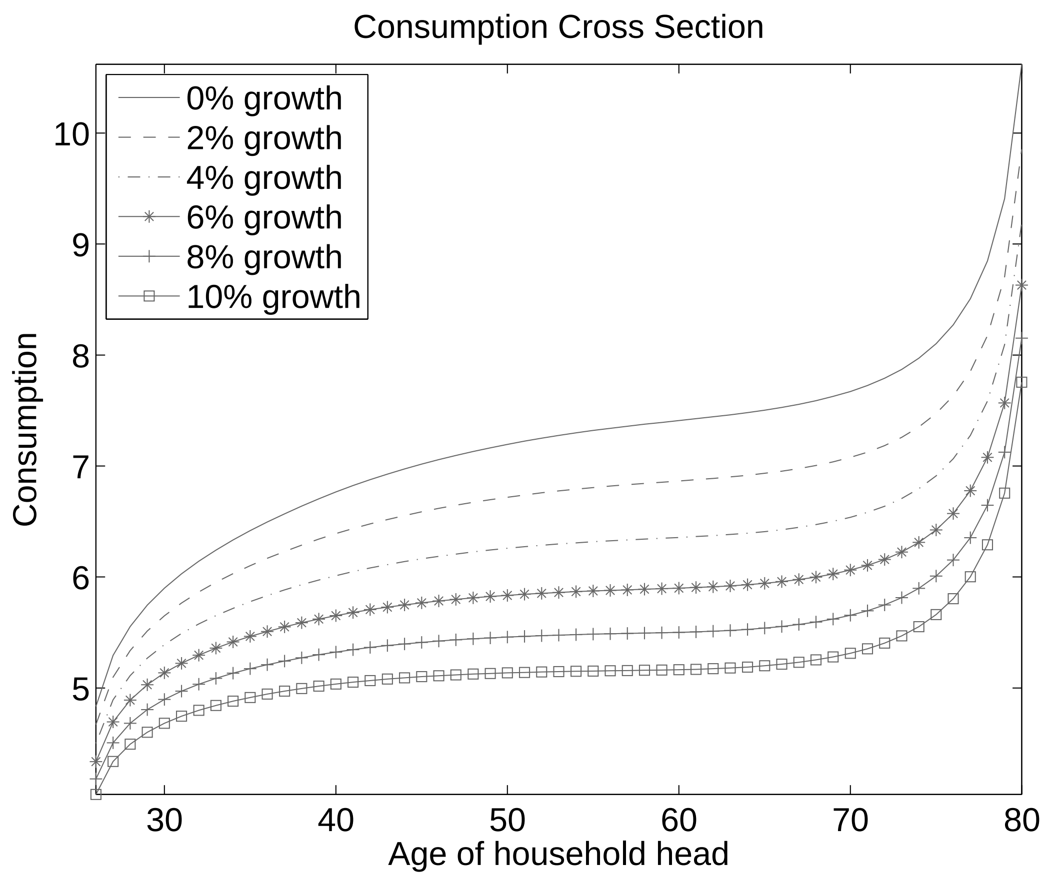

Figure 1.17: Consumption cross-section

1.9 Conclusion

Using Chinese household data, I show that relative wealth plays a crucial role in determining a given household's saving rate. While it is well established in the literature that the rich save a larger share of their income, the fact that it is *relative* wealth that determines the saving rate is a new empirical finding.[23]

The explanation I suggest for this phenomenon is based on the importance of relative consumption and relative asset holdings in households' preferences. Specifically, I create a hybrid utility function model, weighting both the household's consumption and utility from assets with average local consumption and asset holding levels. This

[23]This has been suggested by Duesenberry (1949), but without substantial empirical support.

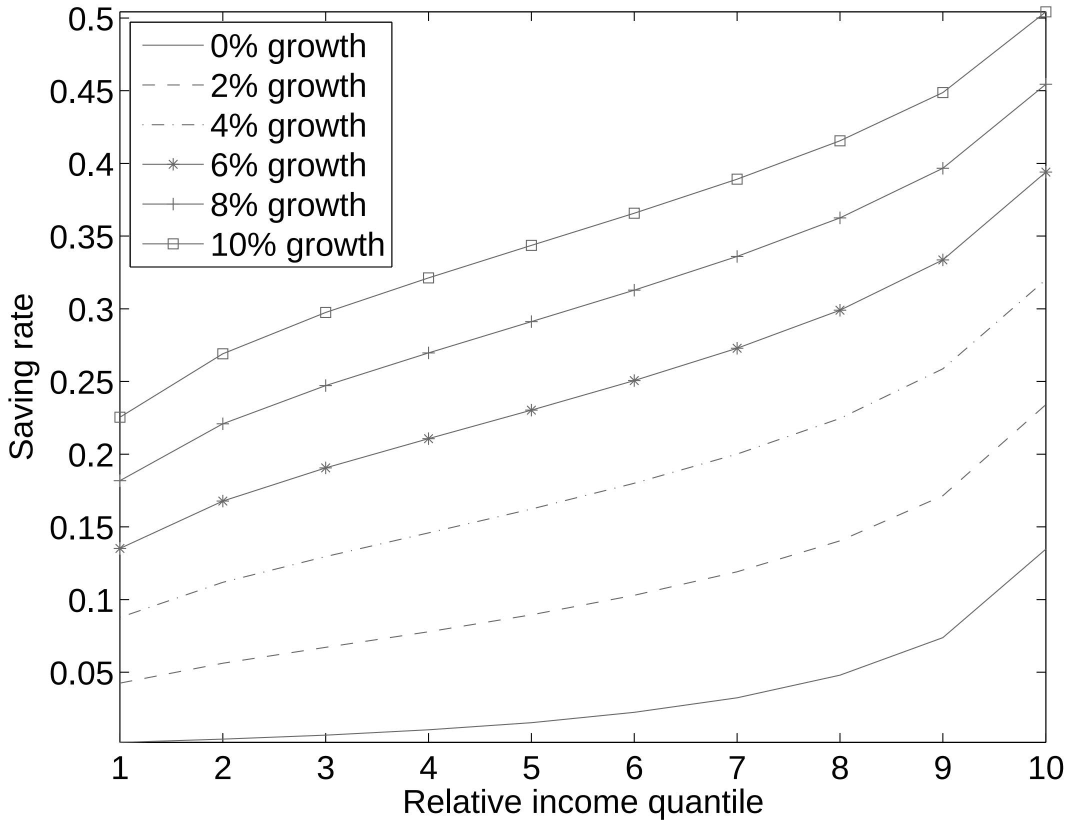

Figure 1.18: Saving rate by local income quantile and aggregate growth rate

utility function is shown, both analytically and via simulation, to recreate behavior evident in the data.

Overall, my findings suggest that social influence may play a large role in determining households' economic behavior, with local rankings in consumption and asset holding possibly functioning as indicators of social success. Such preferences are instrumental in explaining differences in saving behavior not only between rich and poor locations, but also between high and low growth environments. With economic growth, a household may direct more wealth to future use, keeping up with the increase in average income, and commensurately, average consumption and asset holdings. Thus, these preferences, constructed to account for the effect of relative in-

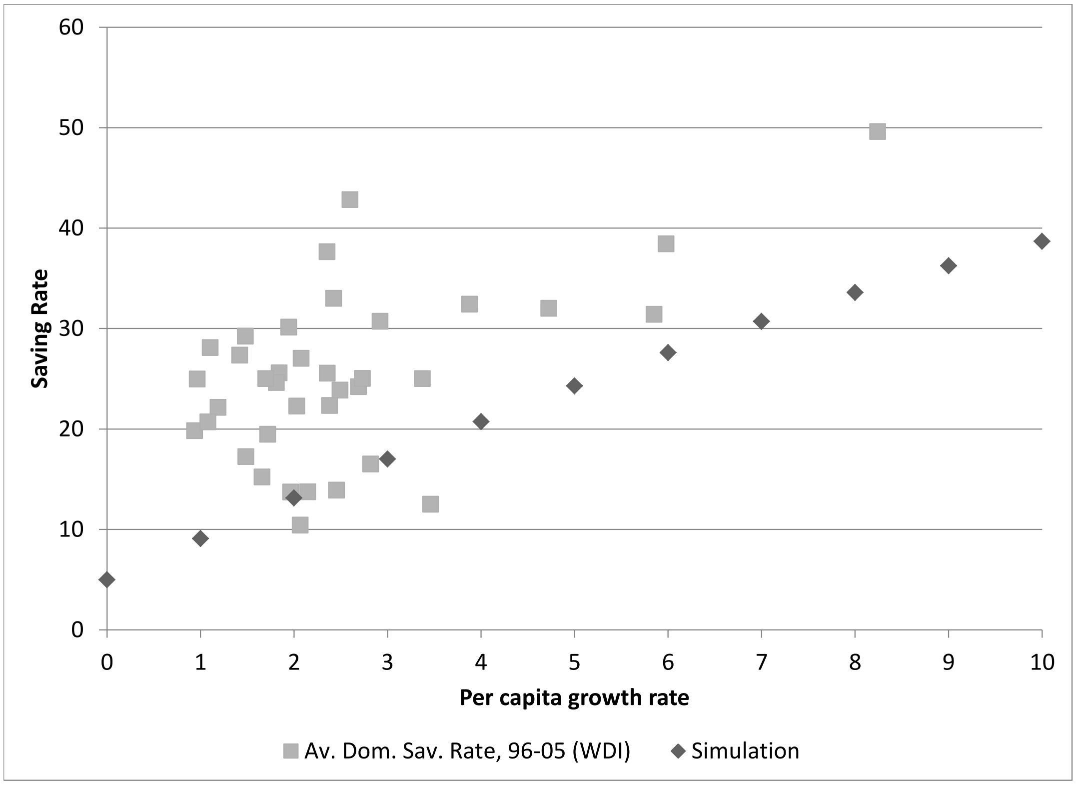

Figure 1.19: Aggregate saving rate by growth, simulation and WDI data

come on household saving rates, are also able to match the empirical effects of rapid growth on saving behavior. I am able to generate age-consumption cross-sections symmetric across growth levels, and simulate a positive relation between aggregate growth and saving rates that closely resembles the data.

From a policy perspective, these findings indicate that the household's saving behavior is strongly influenced by two dimensions of inequality. First, in the spatial dimension, inequality within communities is much more influential than inequality between communities. Second, in the time dimension, inequality between cohorts, as created by growth, also plays a role in the households' saving behavior. A policy aimed at changing saving behavior should therefore focus on these two dimensions of inequality.

Looking forward, the hybrid model has implications for the dynamic behavior of households of different relative wealth levels in response to aggregate and idiosyncratic income shocks. In addition, it provides a framework of analysis for the economic behavior of households in different economic growth, income distribution and social identity environments.

1.A Consumption Shares

City-year income quartile	1	2	3	4	Total
Food	0.556	0.541	0.534	0.516	0.537
Clothing	0.150	0.152	0.142	0.108	0.138
Home equipment and services	0.055	0.063	0.062	0.078	0.065
Medical expense	0.035	0.028	0.027	0.030	0.030
Transportation	0.015	0.018	0.019	0.024	0.019
Communication	0.011	0.016	0.023	0.024	0.018
Entertainment durables	0.013	0.014	0.014	0.024	0.016
Education	0.048	0.051	0.048	0.043	0.048
Culture and recreation	0.016	0.018	0.021	0.029	0.021
Rent and home repairs	0.020	0.020	0.022	0.032	0.023
Utilities	0.050	0.045	0.049	0.046	0.047
Miscellaneous	0.032	0.035	0.040	0.047	0.039

Table 1.8: Consumption Shares in City-Income Quartiles

1.B Proofs

Proposition 1 (Saving rate neutral in equal relative wealth) *If $\{s_t\}_{t=1}^T$ is optimal given $\{y_t\}_{t=1}^T$ and $\{\bar{y}_t\}_{t=1}^T$, then $\forall x > 0$ $\{s_t\}_{t=1}^T$ is optimal for $\{x * y_t\}_{t=1}^T$ and $\{x * \bar{y}_t\}_{t=1}^T$.*

Assumption 1 *Assume $\bar{c}_t = \gamma_1 \bar{y}_t$, $\bar{a}_t = \gamma_2 \bar{y}_t$, so that if $\{\bar{y}_{2,t}\}_{t=1}^T = \{x * \bar{y}_{1,t}\}_{t=1}^T$, then $\{\bar{c}_{2,t}, \bar{a}_{2,t}\}_{t=1}^T = \{x * \bar{c}_{1,t}, x * \bar{a}_{1,t}\}_{t=1}^T$.*

Proof 1 *By Assumption 1, if $\{c_t, a_t\}_{t=1}^T$ is optimal given $\{\bar{y}_t, y_t\}_{t=1}^T$, then $\{x * c_t, x * a_t\}_{t=1}^T$ is optimal given $\{x * \bar{y}_t, x * y_t\}_{t=1}^T$ (where optimality is defined by equations 1.16-1.19), so that the saving rate profile is identical.*

Since the saving rate profile is identical, Assumption 1 holds true. ∎

Lemma 1 (End-of-life asset share increasing in relative wealth) *Given $\{y_t\}_{t=1}^T$, define $Y = \sum_{t=1}^T (1+r)^{T-t} y_t$. If $y_{2,t} = x * y_{1,t} \, \forall t$, where $x > 1$ (household 2 has higher income) and $\bar{y}_2, t = \bar{y}_1, t \, \forall t$ (both households are in a similar average income locality), then $\frac{a_{1,T}}{Y_1} < \frac{a_{2,T}}{Y_2}$ (household 2 will save a larger portion of its lifetime income).*

Proof 2 *By equation 1.17, we get:*

$$\hat{a}_T = \frac{\theta}{\alpha} * \frac{\frac{a_T}{\bar{a}_T} + \chi}{\frac{a_T}{\bar{a}_T}} * \hat{c}_T, \tag{1.20}$$

where hat notations denote a rate of change. Since $\frac{\theta}{\alpha}$, $\frac{\frac{a_T}{\bar{a}_T} + \chi}{\frac{a_T}{\bar{a}_T}} > 0$, $\hat{a}_T > \hat{c}_T$. However, we need to show that $\hat{a}_T > \hat{Y}$. Defining $C = \sum_{t=1}^T (1+r)^{T-t} c_t$, we get $a_T = Y - C$, so we then need to show $\hat{a}_T > \hat{C} = \sum_{t=1}^T (1+r)^{T-t} c_t \hat{c}_t / \sum_{t=1}^T (1+r)^{T-t} c_t$.

Combining equations 1.16 and 1.17, we get:

$$\left(\frac{c_t}{\hat{c}_t}\right)^{-\theta} = \lambda \sum_{j=1}^{T-t+1} \left(\frac{a_{T+1-j}}{\bar{a}_{T+1-j}} + \chi\right)^{-\alpha}. \tag{1.21}$$

Therefore the change rate of consumption at time t is a weighted average of the change rates of asset holdings at periods t to T, multiplied by numbers smaller than unity:

$$\hat{c}_t = \frac{\alpha}{\theta} \frac{\sum_{j=1}^{T-t+1} \left(\frac{a_{T+1-j}}{\bar{a}_{T+1-j}} + \chi \right)^{-\alpha} \frac{\frac{a_{T+1-j}}{\bar{a}_{T+1-j}}}{\frac{a_{T+1-j}}{\bar{a}_{T+1-j}} + \chi} \hat{a}_{T+1-j}}{\sum_{j=1}^{T-t+1} \left(\frac{a_{T+1-j}}{\bar{a}_{T+1-j}} + \chi \right)^{-\alpha}}. \tag{1.22}$$

By equation 1.18 (the budget constraint):

$$y_t \hat{y}_t + (1+r) a_{t-1} \hat{a}_{t-1} = c_t \hat{c}_t + a_t \hat{a}_t. \tag{1.23}$$

By definition (in lemma 1), $\hat{y}_t = x - 1 = \hat{y}\, \forall t$.

(proof by contradiction)

Assume $\hat{a}_T < \hat{y}$. From equation 1.20, $\hat{a}_T > \hat{c}_T$. Define accordingly $1 > \delta_{1,T}$, $\delta_{2,T}$ so that $\hat{a}_T = \delta_{1,T}\hat{y}$ and $\hat{c}_T = \delta_{2,T}\hat{y}$. Then:

$$a_{T-1}\hat{a}_{T-1} = \frac{1}{1+r}(\delta_{1,T} a_T + \delta_{2,T} c_T - y)\hat{y}. \tag{1.24}$$

Since $a_{T-1} > \frac{1}{1+r}(\delta_{1,T} a_T + \delta_{2,T} c_T - y)$, as evident from equation 1.18, then $\hat{a}_{T-1} < \hat{y}$. Therefore, as $\hat{a}_T$, $\hat{a}_{T-1} < \hat{y}$, then, by equation 1.22, $\hat{c}_{T-1} < \hat{y}$.

Define accordingly $1 > \delta_{1,T-1}$, $\delta_{2,T-1}$ so that $\hat{a}_{T-1} = \delta_{1,T-1}\hat{y}$ and $\hat{c}_{T-1} = \delta_{2,T-1}\hat{y}$. Show that $\hat{a}_{T-2}$, $\hat{c}_{T-2} < \hat{y}$ and repeat until $t = 1$.

Therefore, if $\hat{a}_T < \hat{y}$, then $\hat{c}_t < \hat{y}\, \forall t = 1, ..., T$, so that $\hat{C} < \hat{y}$. But since $Y = C + a_T$, it cannot be that $\hat{a}_T$ and $\hat{C}$ are both less than $\hat{y}$.

Thus, we have a contradiction, and it must be that $\hat{a}_T > \hat{y}$. ∎

Proposition 2 (Asset share for all periods increasing in relative wealth) *$\forall x > 1$, if $\{a_t\}_{t=1}^T$ is optimal given $\{y_t\}_{t=1}^T$ and $\{\bar{y}_t\}_{t=1}^T$ and $\{a_t^*\}_{t=1}^T$ is optimal given $\{x * y_t\}_{t=1}^T$ and $\{\bar{y}_t\}_{t=1}^T$, then $\frac{a_t^*}{x * y_t} > \frac{a_t}{y_t} \forall t$.*

Proof 3 *Assume* $\hat{a}_t < \hat{y}$*. Then* $\hat{C}_t > \hat{y}$*, and there exists* $j \leq t$ *such that* $\hat{C}_j$*,* $\hat{c}_j > \hat{y}$*. By equation 1.16, we have the following:*

$$\hat{c}_{t+1} = \hat{c}_t + \frac{\lambda \left(\frac{a_t}{\bar{a}_t} + \chi \right)^{-\alpha}}{\left(\frac{c_{t+1}}{\bar{c}_{t+1}} \right)^{-\theta}} \left(\hat{c}_t - \frac{\alpha}{\theta} \frac{\frac{a_t}{\bar{a}_t}}{\frac{a_t}{\bar{a}_t} + \chi} \hat{a}_t \right). \tag{1.25}$$

Since $\hat{c}_j > \hat{y} > \hat{a}_j$*, then by equation 1.25* $\hat{c}_{j+1} > \hat{c}_j$ *and* $\hat{C}_{j+1} > \hat{y}$*. By induction,* $\hat{C}_s > \hat{y} > \hat{a}_s \, \forall s \geq j$*. But in lemma 1 we showed that* $\hat{a}_T > \hat{y}$*. Therefore, we have a contradiction and* $\hat{a}_t > \hat{y} > \hat{C}_t \, \forall t$*.* ∎

1.C Country List

Country	96-05 per capita growth av.	Av. Dom. Sav. Rate, 05
Argentina	1.42	27.34
Australia	2.49	23.85
Austria	1.84	25.58
Belgium	1.81	24.60
Brazil	0.94	19.81
Canada	2.36	25.54
Chile	2.92	30.71
China	8.24	49.60
Denmark	1.70	25.00
Finland	3.37	24.99
France	1.72	19.46
Germany	1.19	22.13
Greece	3.46	12.49
Hong Kong, China	2.43	33.00
India	4.73	31.99
Indonesia	1.48	29.23
Ireland	5.98	38.43
Italy	1.08	20.68
Japan	0.97	24.95
Korea, Rep.	3.88	32.42
Malaysia	2.60	42.82
Mexico	2.38	22.32
Netherlands	2.08	27.02
New Zealand	2.03	22.26
Norway	2.35	37.64
Pakistan	1.66	15.21
Philippines	2.07	10.42
Portugal	1.96	13.71
South Africa	1.49	17.23
Spain	2.69	24.20
Sweden	2.73	25.01
Switzerland	1.10	28.09
Thailand	1.94	30.13
Turkey	2.82	16.49
United Kingdom	2.45	13.89
United States	2.15	13.73
Vietnam	5.85	31.39

Figure 1.20: International Growth and Saving Rates, Source: World Development Indicators, World Bank

CHAPTER II

Growth and saving in an open economy structural change life cycle model – an investigation

2.1 Introduction

The phenomenon of high saving rates accompanying rapid growth is well known in the economic literature. Modigliani (1970) notes it with relation to post world-war II Europe. More recently it is usually identified with the rapidly growing Asian economies. Deaton and Paxson (2000) state, based on cross-country evidence, that a one percent increase in per-capita growth is accompanied by a two percent increase in the saving rate. The data presented in Appendix II seems to be in accordance with that estimate.

However, the standard infinite horizon representative agent models predict the opposite - the higher is the expected growth, the more is the agent to borrow. In the international context, this means that a relatively poor but rapidly developing country is supposed to run trade deficits, essentially borrowing from rich countries at the expense of its future wealth. The opposite is often observed in reality.

Initially it was suggested that it is the higher propensity to save that is causing the higher growth, a-là a Solow style model. Later it became apparent that it is higher growth that precedes higher saving rates and not the other way around;[1] hence it is more likely that causality runs from growth to saving. While a positive causality running between saving and growth cannot be ruled out, clearly the saving decision must be optimal given growth expectations.[2]

Modigliani (1966) attempts to explain the positive causality from growth to saving using a life-cycle model. The intuition behind Modigliani's explanation is simple:

[1] See Carroll and Weil (1994) and Andersson (1999).

[2] For example, Aghion et al. (2006) suggest a model in which domestic savings allow domestic banks to co-finance foreign technologically advanced investment and thus foster growth.

since it is the young that save (toward retirement) and the old that dissave (being retired), growth, making the young richer compared to the old, should raise the saving rate. However, as Tobin (1967) and Carroll and Summers (1991) note, this explanation requires that growth increases income from cohort (age group) to cohort, but not within a cohort (i.e. growth does not affect an existing agent's lifetime income, but rather increases the lifetime income of a "newborn" agent with respect to the previous-generation agent), otherwise the young, expecting higher future income, would lose their motivation to save. Carroll and Summers (1991) bring evidence to refute this assumption (i.e. they show that growth increases the income of people of all ages), and argue that this invalidates Modigliani's life-cycle explanation completely. Another attempt to explain this direction of causality is made by Carroll and Weil (1994) and Carroll et al. (2000), using a habit formation mechanism. The latter paper also contains an illuminating literature review of the topic. While habit formation may indeed create a positive relation from growth to saving, by shifting consumption to the future (where the agent is more "spoiled"), it is far from a standard feature of the utility function, and examining Chinese data, Chamon and Prasad (2010) find no evidence of it.

In addition to its effect on the saving rate, rapid growth seems to have two more closely related effects: In Carroll and Summers (1991), the authors also find that households in rapidly growing countries experience a higher rate of increase in their consumption expenditure compared to their peers in lower growth countries. This phenomenon, also observed in Lee et al. (2006), which Carroll et al. (2000) seek to explain with habit formation, is also inconsistent with the standard PIH and life cycle models. A second related phenomenon, observed in Carroll and Summers (1991), is that the cross-sectional profile of consumption by age is quite similar across countries with very different aggregate growth rates.

This paper attempts to quantify and measure the ability of a structural change life-cycle model to generate the stylized facts, namely a positive relation from aggregate growth to the aggregate saving ratio (with growth fully affecting individual cohorts), correlation between the aggregate growth rate and the household consumption expenditure growth rate over the life cycle, and consumption-age cross sectional distribution that is largely invariant to aggregate growth. The term "structural change" is used here to refer to changes in the division between traditional (with slower growth in technology) and modern (with faster growth in technology) production sectors and its effect on the relative price level. In essence, if growth, increasing the household's future income, can be said to be creating a "life cycle income effect" leading to borrowing, the price level change caused by structural change can be said to be creating a "life cycle substitution effect", diverting expenditure toward the later years of the household and thus encouraging saving. The sum of the two effects decides the direction of the correlation between growth and saving.

In addition, the paper seeks to ascertain whether further augmenting the life-cycle model in a "reasonable" manner can substantially strengthen its ability to generate the stylized facts. Of the several possibly relevant additions, this paper focuses on incorporating the effects of uncertainty about future growth and credit constraints. The effect of each element on the savings to GDP ratio in the life cycle setting is examined both separately and in conjunction with the other elements.

Section two presents the life-cycle model used. The structural change two-sector production model is in section three. Sections four and five respectively describe the modeling of growth uncertainty and credit constraints. The respective foci of sections six and seven is the calibration and simulation of the model. Section eight concludes.

2.2 The Life-cycle Model

The demand side of the economy is represented by a household life-cycle model. Each period t one household is created and one is destroyed (i.e. the age distribution is uniform). Each household exists for J periods, so that there exists one household for each household age j from one to J. Each household has the following identical lifetime utility function:

$$U = \sum_{j=1}^{J} \beta^{j-1} u\left(j, c(j,t,g,r)\right) \tag{2.1}$$

where g is the growth rate in the economy, and is assumed, for now, to be constant and certain. $c(j,t,g,r)$ is thus the consumption of a household of age j at time t, given growth rate g and a constant international interest rate r.[3]

The household's within-period utility varies with age, and is of the following form:

$$u\left(j, c(j,t,g,r)\right) = q\left(j\right) \frac{c(j,t,g,r)^{1-\theta} - 1}{1-\theta} \tag{2.2}$$

where $q(j)$ is what shall be referred to as the "life-cycle multiplier", putting weights

[3] The small open economy assumption seems appropriate for countries starting out on the development path. It may become less appropriate as the economy grows, as it may reach a size sufficient to influence international capital markets. This paper abstracts away from such situations in the interest of simplicity.

on utility in different periods in the household's life cycle. This is an element that is often thought of as the average number of equivalent adults for household age j.[4] I will refer to the path of the life-cycle multiplier over the life of the household as the "household regime" (i.e. $\{q(j)\}_{j=1}^{J}$).

One could argue that since we observe the average number of children per household diminishes with economic modernization, the household regime should vary with the level of development. However, there is much evidence that as the number of children per household falls, the expenditure share spent on each child rises.[5]

Another issue is the effect of development on the norm, common in many developing countries, of elderly parents moving in with their children. There is firm evidence[6] that prolonged development, of the kind this paper is focused on, lowers the rate in which this phenomenon takes place. Thus a young person expecting long term growth cannot rely on support from his or her children after retirement. Therefore, based on the understanding that the household regime of developing countries converges to that of developed countries with growth, I choose to abstract away from this element.

In light of the above, and following evidence (discussed in part six of the paper) that the consumption over age profile in the US is very stable, I assume the household regime to be time invariant.

The household's effective labor units, l, are assumed to change over the household's age, j. This is attributed to changes in experience, improved position in the workplace, as well as physical and mental abilities varying with age. I assume the total number of effective units of labor at the economy to be constant at unity:

$$\sum_{j=1}^{J} l(j) = 1 \tag{2.3}$$

The household's labor income over the life-cycle is therefore dependent upon household age j and the prevailing wage per effective unit of labor at the time period, $W(t,g)$:

$$w(j,t,g) = l(j)W(t,g) \tag{2.4}$$

where $w(j,t,g)$ is the labor income of a household of age j at time t under a prevailing growth rate of g. As will be shown in part three, $W(t,g)$ is growing at rate g, hence

[4] See, for example, Tobin (1967), Modigliani (1970, 1986), Mariger (1987) and Gourinchas and Parker (2002).

[5] A seminal paper in this "quantity vs. quality" literature is Becker and Lewis (1973).

[6] See UN (2005).

growth fully affects the household's labor income over the life cycle (i.e. the "life cycle income effect"). Thus this model is not subject to the Carroll and Summers (1991) critique.

Under the assumption that the economy is small and open to capital flows (i.e. the interest rate, r, is constant), discounted lifetime income is:

$$\Omega\left(t,g,r\right)=\sum_{j=1}^{J}\left(1+r\right)^{-j+1}w\left(j,t+j-1,g\right)=\sum_{j=1}^{J}\left(1+r\right)^{-j+1}l(j)W(t+j-1,g) \tag{2.5}$$

As specified in part three, the price level P is a function of the growth rate and time. The discounted lifetime consumption expenditure of a household created at time t in an economy with a constant growth rate of g is therefore:

$$C\left(t,g,r\right)=\sum_{j=1}^{J}\left(1+r\right)^{-j+1}P\left(t+j-1,g\right)c\left(j,t+j-1,g,r\right) \tag{2.6}$$

Each individual household is thus facing the following utility maximization problem at the time of its creation:

$$\begin{aligned}&\max\sum_{j=1}^{J}\beta^{j-1}q\left(j\right)\frac{c(j,t+j-1,g,r)^{1-\theta}-1}{1-\theta}\\&s.t.\\&C(t,g,r)=\Omega\left(t,g,r\right)\end{aligned} \tag{2.7}$$

Where the households simply optimize their respective consumption streams under the constraint of their lifetime labor income.

2.3 Structural Change - The Two Sector Supply Side

The purpose of two sector setup of the supply side in this paper is to model unbalanced growth, i.e. the structural change that occurs as the economy grows, and the resulting shift in the relative prices. As one sector's productivity grows faster than the other's, factor mobility causes the goods produced by the traditional (slower productivity growth) sector to increase in price relative to modern (faster productivity growth) goods. Assuming that the two kinds of goods are not perfect substitutes, the result is an increase in the overall price level in the economy. This result, in different formulations, can separately be credited to Baumol (1967), Balassa (1964) and Samuelson (1964). The first generally refers to a structural changes in a single growing economy, while the latter two use the same principal to explain cross-country differences.

The production side modeling in this paper has been influenced by the Obstfeld and Rogoff (1996) formulation of the Balassa-Samuelson hypothesis, as well as by Ngai and Pissarides (2007).

The phenomenon modeled here is that of developing countries, catching up with the world leaders through the gradual adoption of advanced organizations and technology. This catching-up is thus driven by increased productivity (as opposed to just capital deepening), and is marked by high rates of growth spanning a relatively long period of time. This increase in productivity is assumed to manifest itself in some sectors of the economy more than in others, changing the relative prices and consumption shares of the sectors. Chamon and Prasad (2010), for example, document a large change in the composition of household consumption in China, and largely credit that change with the increase in household saving rates.

The model comprises two sectors of production: a modern sector, which enjoys rapid productivity growth (through technological advancement), and a traditional sector, which is slow to take advantage of technological improvements (if at all). The modern sector's product serves as the numeraire. In addition, capital is assumed to be produced by the modern sector only.

Due to the open economy assumption of this model, the goods produced by at least one of the sectors in question must be non-tradable to a large degree. Otherwise, the price ratio between the sectors would be set internationally irrespective of their relative productivity in the domestic economy. For ease of calibration, I identify the traditional sector to be the largely non-tradable service sector.[7] Notice however that this need not be exclusive – other sectors (e.g. housing) which are largely non-tradable and less-affected by technology would fall under the 'traditional' definition in this model.

I assume for simplicity that production is Cobb-Douglas in both sectors, and that factor intensity (α) is identical.[8]

[7] In this I take after much of the relevant literature, from Baumol (1967) to Ngai and Pissarides (2007). A paper directly investigating the link between growth and service sector size is Kravis and Summers (1983). Lee and Wolpin (2006) also pins the cause of the growth in US service sector employment on technological change (as well as product and capital prices).

[8] In this I follow Ngai and Pissarides (2007). While it may be quite reasonable to expect a higher capital intensity in the modern sector, such specification should not alter the results qualitatively. In essence, the increase in the relative prices of traditional goods with technological advancement would be amplified by capital deepening.

$$Y_{M,t} = A_{M,t} K_{M,t}^{\alpha} L_{M,t}^{1-\alpha} \tag{2.8}$$

$$Y_{T,t} = A_{T,t} K_{T,t}^{\alpha} L_{T,t}^{1-\alpha} \tag{2.9}$$

where the growth rates of productivities in the two sectors are assumed constant:

$$A_{M,t} = (1 + g_M)^t \tag{2.10}$$

$$A_{T,t} = (1 + g_T)^t \tag{2.11}$$

Assuming free factor movement between the two sectors, the labor wage and the rental rate of capital would be equal.[9] The result is an equal capital to labor ratio between the sectors:

$$\frac{K_{M,t}}{L_{M,t}} = \frac{K_{T,t}}{L_{T,t}} = \frac{K_t}{L} = K_t \tag{2.12}$$

where $K_t = K_{M,t} + K_{T,t}$ and $L = L_{M,t} + L_{T,t} = 1$.

Assuming an open economy, we have that the rental rate of capital, R (which equals the interest rate, r, plus the depreciation rate, δ), remains constant, so that:

$$K_t = \left[\frac{\alpha A_{M,t}}{R}\right]^{\frac{1}{1-\alpha}} = (1 + g_M)^{\frac{t}{1-\alpha}} \left[\frac{\alpha}{R}\right]^{\frac{1}{1-\alpha}} \tag{2.13}$$

$$\Delta K_t \approx K_t \frac{g_M}{1-\alpha} \tag{2.14}$$

where the latter equation is using the continuous time change to approximate the discrete time change.

The relative price of the traditional good in terms of the modern good is thus dependent on the relation between the two levels of productivity:

$$p_t = \frac{A_{M,t}}{A_{T,t}} \equiv (1 + \phi)^t \tag{2.15}$$

where $1 + \phi \equiv \frac{1+g_M}{1+g_T}$ and $\phi \approx g_M - g_T$ (a continuous time approximation).

In this paper I will assume that the driving force behind rapid economic growth

[9] Here is where the assumption that capital is a modern good comes into play. Otherwise one could invest in traditional goods and profit from their appreciation in terms of modern goods. For some traditional goods (e.g. services) this assumption makes sense, and for others (e.g. real estate) it does not, however in the sake of simplicity, and since it does not affect results significantly, I abstract away from investment in traditional goods.

is productivity improvements in the modern sector, and that productivity in the traditional sector is mostly stagnant. Accordingly, g_T is assumed to be zero, so that ϕ is equal to g_M. Both aggregate growth and relative prices are therefore entirely dependent upon productivity growth in the modern sector.

Output would therefore be:

$$Y_t = Y_{M,t} + p_t Y_{T,t} = A_{M,t} K_t^{\alpha} \tag{2.16}$$

With constant capital rental rate ($R = r + \delta$) the above equation can be rewritten as:

$$Y_t = (1 + g_M)^{\frac{t}{1-\alpha}} \left[\frac{\alpha}{R}\right]^{\frac{\alpha}{1-\alpha}} \tag{2.17}$$

Clearly g, the growth rate of output, is approximately $\frac{g_M}{1-\alpha}$. The above equation can therefore be rewritten as:

$$Y(t, g, r) = (1 + g)^t \left[\frac{\alpha}{R}\right]^{\frac{\alpha}{1-\alpha}} \tag{2.18}$$

Since in this model the wage per effective labor unit is a constant fraction of output ($W(t, g) = (1 - \alpha)Y(t, g)$) then clearly it increases at rate g, as stated in the previous section.

It also follows that, in equilibrium, the necessary investment to output ratio under a stable growth path would approximately be:

$$i(g, r) \equiv \frac{\Delta K_t + \delta K_t}{Y_t} \approx \frac{\alpha}{R}\left(\frac{g_M}{1-\alpha} + \delta\right) \approx \frac{\alpha}{R}(g + \delta) \tag{2.19}$$

Next I assume a CES aggregation function, through which the two types of goods form together a composite good:

$$C_t = \left[\gamma^{\frac{1}{\epsilon}} C_{M,t}^{\frac{\epsilon-1}{\epsilon}} + (1-\gamma)^{\frac{1}{\epsilon}} C_{T,t}^{\frac{\epsilon-1}{\epsilon}}\right]^{\frac{\epsilon}{\epsilon-1}} \tag{2.20}$$

Notice that a γ value of one corresponds to the standard one sector case whereas a γ value of zero corresponds to a completely stagnant economy. ϵ is the elasticity of substitution between the two goods. An ϵ of zero is tantamount to the Baumolian case of constant shares between the two goods (i.e. a Leontief function). An ϵ of one is equivalent to using a Cobb-Douglas aggregation function, where the expenditure shares are constant. A high enough ϵ will make the two kinds of goods very substitutable, leading to modern goods largely replacing traditional goods.

The share of the traditional sector is a function of p the relative price of the traditional good in terms of the modern good:

$$S_t^T = \frac{1}{1+\frac{\gamma}{1-\gamma}p_t^{\epsilon-1}} = \frac{1}{1+\frac{\gamma}{1-\gamma}(1+\phi)^{t(\epsilon-1)}} \tag{2.21}$$

The aggregate price level in the economy, reflecting relative prices and shares of the two kinds of goods, is the minimal cost, in terms of modern goods, of one unit of composite good:

$$P_t = \left[\gamma + (1-\gamma)p_t^{1-\epsilon}\right]^{\frac{1}{1-\epsilon}} = \left[\gamma + (1-\gamma)(1+\phi)^{t(1-\epsilon)}\right]^{\frac{1}{1-\epsilon}} \tag{2.22}$$

Hence $\hat{P}_t$, the growth rate of the price level, is generally not constant over time (save for the cases where γ is either zero or one or the unique Cobb-Douglas case where $\epsilon = 1$).

$$\hat{P}_t = \frac{(1-\gamma)p_t^{1-\epsilon}}{\gamma + (1-\gamma)p_t^{1-\epsilon}}\phi = \frac{1}{1+\frac{\gamma}{1-\gamma}(1+\phi)^{t(\epsilon-1)}}\phi = S_t^T\phi \tag{2.23}$$

Thus, for an ϵ between zero and one will have $\hat{P}_t$ converge toward ϕ, and an ϵ above one will have it converge to zero. For the case where $\epsilon = 1$ we get that $\hat{P}_t$ is constant at $(1-\gamma)\phi$.

Notice that this increase in the price level is not the same as inflation, in the sense that it denotes an actual decrease in purchasing power, whereas inflation usually translates to increase in both nominal price and income, leaving actual purchasing power unchanged over time.

Notice also that for the case of ϵ between zero and one, even as the share of the traditional sector converges to one, output growth, at $\frac{g_M}{1-\alpha}$, is still markedly faster than price growth, which asymptotes toward g_M. Therefore, while the growth of effective purchasing power in the economy slows, it does not grind to a halt, as in Baumol (1967). This is a result of having more than one factor of production in the model, i.e. enabling capital deepening to enhance productivity growth.

With a γ under one, we have a price level growth rate that is positively correlated with the growth rate of the economy. Thus a household in a high growth regime can not only expect its income to increase over time, but also its cost of living. Since in later years (i.e. in retirement) its income is bound to drop but the cost of living is expected to keep increasing, maintaining the desired level of consumption would require more savings. This is the aforementioned "life cycle substitution effect", causing households to substitute between current and future consumption.

2.3.1 Choice of Numeraire

As detailed above, output (Y) is in units of the modern good. It may be argued that a more correct way to measure output would be to adjust for the price level, i.e. that output should equal $\frac{Y}{P}$. This is a critical measurement issue, as doing so would nullify the effect of structural change. The foremost question one should therefore ask in this context is: to what extent do official price indices take account of structural change effects?

Deaton (2005) argues that the failure to properly measure nonexchanged goods and services (mostly imputed housing cost and home production) causes systematic discrepancies between income measured by national accounts and income measured by household surveys, and that the former systematically overstate growth.

In addition, there are a number of other factors that are likely to cause an under measurement of the effect of structural change on prices. First, government services are notoriously difficult to price. For example, is an increase in the expense on public education inflationary, or is it driven by an increase in the quality of teaching? Second, the price of housing, also mentioned by Deaton (2005), which would fall under the traditional good category (construction technology is fairly stagnant and houses are largely non-tradable), is often omitted from the price index, or, as is the case in the US, is substituted for by imputed rent. As the price of housing is rising much faster than renting cost, it is under-represented in the price index.[10] Third, for the effect we are looking for, it is the price level faced by the households that is of importance. Forming an index of US household consumption (a weighted average of goods, services and government expenditure prices), I find that it is rising faster than the GDP deflater. Over years 1960-2008 this consumption index rose by about 25% faster than the GDP deflater.

2.4 Uncertainty of Future Growth

Up till now I have assumed that households know with certainty that growth will continue at a constant rate throughout their lifetime. While, barring a major catastrophe, the fluctuation in lifetime income an average household in the developed world can expect are relatively mild, that may not be the case for a household in a rapidly developing country.[11] Such a country is in a process of catching up to

[10] The HPI index, using actual house purchases, rose about 50% over years 2000-2007, more than twice than the rent imputed housing cost used in the CPI, which rose by about 24%.

[11] This model has no idiosyncratic shocks. Such shocks, assuming that they are not influenced by the rate of growth, are unlikely to explain the higher rate of saving in a rapidly growing economy.

the developed world "frontier", and is able to grow rapidly by adopting existing technologies and practices. It is, however, bound to run out of such technologies once it reaches the development frontier. Its growth rate should then slow down and be in line with the rest of the developed countries. Such a stop can be abrupt (e.g. Japan) or slow and gradual (e.g. Hong-Kong and Singapore). It is in fact very difficult to establish the point in which a country fulfills its economic potential and reaches the frontier.

In addition, it is also possible for a country to drop out of the catching-up process altogether (without reaching the development frontier). After all, as Parente and Prescott (1993) note, the vast majority of developing countries are not in the process of catching up. Being the exception rather than the rule, one can certainly imagine that this process may stop or be derailed; that the unique circumstances causing it may change, or put differently, that whatever factor that was holding the country back to begin with may return. The catching up process might depend on a unique political situation, on geopolitical stability, on the state of the world economy or even on natural events (rainfall, flooding etc.). The effect of the 1997 Asian crisis on some of the local economies is a possible example of such a case, as Indonesia, Malaysia, the Philippines and Thailand have yet to regain the levels of growth enjoyed pre-crisis.

While we do see several countries that are enjoying or have enjoyed a fairly stable stretch of rapid growth over several decades, it may be that their success was far from certain ex-ante. For example, a Japanese household living during the second half of the 20th century had no way of knowing how long would rapid growth continue, and had to take the eventuality of an abrupt stop at any point in time into consideration when making its saving decisions.

Given all of the above, it seems fit to introduce uncertainty into the model. This is done in a fairly simple fashion: Each period, all households perceive a constant probability λ for the country to drop out of its high growth regime (where it grows at rate g_H) into an absorbing low growth regime (where growth equals g_L).[12] This drop is considered to be permanent by all existing households (the country will grow slowly for the rest of their lives). This can be formulated in the form of a Bellman equation:

[12] It is reasonable to believe that λ should actually become higher as the country approaches the development frontier. I abstract away from such variations in λ for simplicity.

$$V^H(a,j) = \max_c \quad u(c,j) + \beta \left\{ \lambda V^L(a',j+1) + (1-\lambda)V^H(a',j+1) \right\} \tag{2.24}$$

$$s.t. \quad a' = (a + w - cP)(1+r), \quad a_{J+1} = 0 \tag{2.25}$$

Where a is asset stock, j is household age, and:

$$V^L(a,j) = \max_{c_t} \quad \sum_{t=j}^{J} \beta^{t-j} u(c_t, t)$$

$$s.t. \quad \sum_{t=j}^{J} (1+r)^{j-t} c_t P_t = a + \sum_{t=j}^{J} (1+r)^{j-t} l(t) W(j) (1+g_L)^{t-j}$$

However, since it is the rapid growth scenario that is of interest here, the households pleasantly discover each period that the high growth regime is still in effect. Thus they re-optimize their consumption decisions each period. The effect of this form of uncertainty is that lifetime wealth is gradually realized to be higher than expected over the household lifetime, and the consumption profile is constantly updated upwards. This effect is obviously stronger the higher is the high regime's growth rate compared to that of the low regime.

Equations (23) and (24) lead to the following Euler equation:

$$u_c(c_j^H, j) = (1+r)\beta \left\{ \frac{\lambda u_c(c_{j+1}^L, j+1)}{(1+\phi_L)^{1-\gamma}} + \frac{(1-\lambda) u_c(c_{j+1}^H, j+1)}{(1+\phi_H)^{1-\gamma}} \right\} \tag{2.26}$$

Where c_{j+1}^H is the optimal consumption choice at period $j+1$ given that the high growth regime remains in effect, and c_{j+1}^L is optimal consumption given a regime switch to low growth at period $j+1$.

2.5 Credit Constraints

Rapid growth creates a great disparity in terms of lifetime wealth between younger and older cohorts in the economy. This disparity is so substantial that the saving decisions of older generations become negligible, and aggregate saving behavior is determined almost entirely by younger generations. Modigliani (1970) argues that this disparity should increase the aggregate saving rate, as the young save (for retirement) and the old dissave (during retirement). Carroll and Summers (1991) argue that growth affects the future income of the younger cohorts, therefore causing them to save less and even dissave at the expense of future income, reducing the aggregate saving rate (i.e. what i refer to as the "life cycle income effect").

Paramount to the latter is the ability of the younger cohorts to take advantage of their high future income through borrowing. Such borrowing, however, would put these cohorts in the realm of negative asset value. This raises the issue of incentives to default on debt through bankruptcy for the individual household (when the cost of bankruptcy is lower than the debt itself), and the possibility of aggregate risk to lending institutions under growth uncertainty - if the entire economy switches into a low growth regime, a large portion of indebted households may find their future income much lower compared to their current debt, and default, bankrupting the lending institutions in the process.

A possible implication of these issues is a credit constraint, preventing the households from borrowing too much at the expense of the future, and restricting the dissaving of the young cohorts. As Jappelli and Pagano (1994) show, using a three period overlapping generations model, a credit constraint may raise the saving rate and strengthen the effect of growth on saving.

In this paper I use the following credit constraint formulation:

$$a_j \geq \kappa \Omega_j \tag{2.27}$$

where Ω_j is expected discounted future labor income, i.e.

$$\Omega_j = \sum_{t=j}^{J} (1+r)^{j-t} E_j \left[w_t \right] \tag{2.28}$$

and κ is some number between 0 (a non-negativity constraint) and -1 (no constraint).

The value of κ may reflect the strength and accessibility of the financial system, as well intergenerational lending. Since the ability of individuals to amass debt is very limited even in developed economies, one would expect a κ near zero to be a reasonable approximation. A possible caveat, however, is that intergenerational transfers (such as bequests) could move young households away from the credit constraints and allow a dissaving pattern similar to the unconstrained case.

I use the algorithm described in Mariger (1987) to compute consumption paths given credit constraints.

2.6 Calibration

In order to examine the actual impact of adding the previously described elements on the ability of the life cycle model to explain the positive correlation of growth and saving, the parameters of the model need to first be calibrated. Of these parameters, some have fairly accepted values (or ranges) in the economic literature (e.g. α, β and θ), while others are more elusive (e.g. $q(j)$ and $l(j)$). For the former, I simply use values within the common ranges, as follows:

$$\alpha = \frac{1}{3}, \quad \beta = 0.95, \quad \delta = 0.05, \quad \theta = 4$$

where α is the exponent of capital in the production function, β is the time discounting parameter, δ is the annual depreciation rate and θ is the risk averseness parameter (the inverse of the intertemporal elasticity of substitution).

Additional parameters that require calibration for the structural change part of the model are ϵ (the elasticity of substitution between the traditional and modern good), γ and p_t. As equation (21) shows, a lot of information on these variables could be retrieved from the share of the traditional good (S^T) over time, could such a series be observed. Specifically, both ϵ and $\frac{\gamma}{1-\gamma}p_t$ can be backed out. [13]

$$\frac{\gamma}{1-\gamma} p_t^{\epsilon-1} = \frac{1 - S_t^T}{S_t^T} \tag{2.29}$$

$$\epsilon \approx \ln\left(\frac{(1 - S_{t_1}^T)S_{t_0}^T}{(1 - S_{t_0}^T)S_{t_1}^T}\right)\frac{1}{(t_1 - t_0)\phi} + 1 \tag{2.30}$$

where t_0 is the first period in the series and t_1 the last.

$$p_t = \left(\frac{1 - S_t^T}{S_t^T}\frac{1-\gamma}{\gamma}\right)^{\frac{1}{\epsilon-1}} \tag{2.31}$$

In order to calibrate the structural change parameters, I choose to identify the traditional sector with the service sector. The service sector is largely non-tradable, and is often less affected by technological progress than the manufacturing sector (though, admittedly, counterexamples can be found). The service sector share is also often found to be correlated with the level of economic development. [14]

I use BEA data for the composition of US consumption, years 1960-2008, to cali-

[13] While γ cannot be identified from such a series, such an identification is not necessary for this exercise.

[14] See discussion in part three.

brate ϵ. Computing the share of service consumption (defined here as personal services consumption plus total government expenditure over total personal and government consumption), and using equation (30), I get an ϵ value of about 0.065, reflecting very little substitutability between goods and services. This is actually very close to the Baumolian case of zero substitutability.

In addition, the life cycle parameters need to be calibrated. My calibration strategy for $q(j)$ and $l(j)$ is detailed below.

Time periods are measured in years. Households are assumed to exist for 55 years (approximately ages 25 to 80).[15]

Using the fact that $c(j,t,g,r) = c(j,(t-j+1)+j-1,g,r)$, equation (36) from appendix I can be reformulated, and $q(j)$, the household's life-cycle multiplier, can be backed out of actual household consumption data, assuming that $1+r \approx \frac{1}{\beta}$. [16] First, the equation is rewritten in matrix form:

$$C_t = X \cdot q/(Y_t q) \tag{2.32}$$

where '$\cdot$' denotes pointwise (element by element, or Hadamard) multiplication and '/' denotes pointwise division.

$C_t(j)$ is the consumption share of cohort j at period t

$X(j) = (1+g)^{-j+1}$

$Y_t(i,j) = (1+r)^{-j+1} P(t-i+j)^{\frac{\theta-1}{\theta}}$

In other words, the model attributes consumption variation across cohorts to the different values of $q(j)$ as well as growth and price level change.[17] Equation (31) above can be rewritten in the following form:

$$C_t / X\iota' \cdot Y_t q = q \tag{2.33}$$

where ι is a $J \times 1$ vector of ones.

Thus q is simply the first eigenvector of the matrix $C_t/X\iota \cdot Y_t$.

[15]This assumption may seem somewhat problematic, considering the changes in longevity and social structure that come hand in hand with economic development. However, I believe it to still be a good approximation: while life expectancy is lower in less developed countries, the differences conditional on fact that the individual has reached adulthood and was able to create an independent household are smaller.

[16]This assumption seems plausible in a developed country settings, considering capital market sophistication and openness.

[17]It is worth noting that other elements, (i.e. growth uncertainty and credit constraints), can cause and/or influence the cohort consumption distribution. Since the calibration is using US data for periods averaging about 2% annual per capita growth, then as will be demonstrated later on, these elements should not have a substantial influence.

I turn to US Consumer Expenditure (CE) survey data on household consumption, years 1984-2005.[18] As can be seen in figure 1, consumption patterns in the data remain very consistent over the years.[19] Normalizing the data from different years to get consumption shares, averaging between the years, and using a cubic spline extrapolation (raised to the power of θ), gives an approximation of $q(j)$, which is then normalized to sum up to unity.

The resulting household regime is hump shaped, peaking when the age of the household head is around 50 (this corresponds to household age 25). This estimation is using only US data, and cannot be claimed to be an accurate representation of life cycle consumption preferences elsewhere. However, its hump shaped result has a general ring of truth to it: we would generally expect a household to require relatively little consumption starting out, as it is composed of just two young individuals. As the household expands, so do its needs - a larger family requires more food, clothing and living space. Consumption generally peaks as the younger generation is preparing, through education and/or acquiring initial property, to establish its own independent household. Later on, as the children leave, the parents can cut down on consumption. As the household draws near to its end, a possible death of a spouse can further decrease consumption. While such regimes are most likely not identical between countries, they nonetheless seem to be broadly similar.

Having experimented with a variety of household regimes, I find that, save for extreme and unreasonable cases, using different household regimes of the general bell shape described above does not affect results substantially.

Calibrating $l(j)$, the effective labor units of the household over its life-cycle, is even more straight forward: I again use US Consumer Expenditure survey data (years 1984-2005) on the wage income of different household cohorts to extrapolate the household's lifetime wage path. The calibration is presented in figure 2.

The hump shape of the estimated effective labor units over the life cycle is not surprising: effective labor units increase with age, likely due to experience and acquired position, peaking, much like the household regime, around the household head age of 50. They then start to fall rapidly, likely due to retirement and natural physical deterioration.

Notice that this formulation implicitly implies that the distribution of effective

[18] According to the World Bank's World Development Indicators, the average US real interest rate was just over 5% in this period, making the approximation of $(1+r)\beta \approx 1$ very good for a β value of 0.95.

[19] This is in line with the findings of Kotlikoff and Summers (1981) and Ando and Kennickell (1986) with respect to the US and Japan.

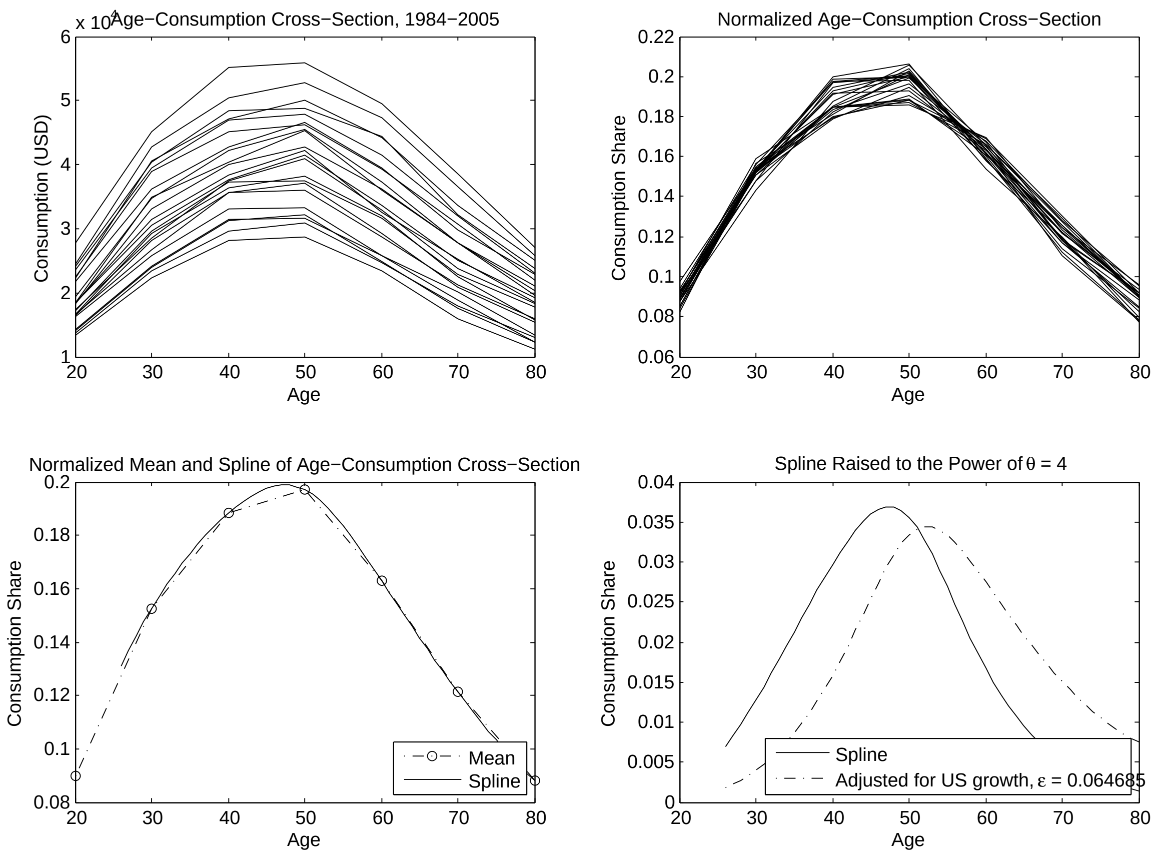

Figure 2.1: Calibrating the household regime, the weighting (q) on the household consumption utility by period

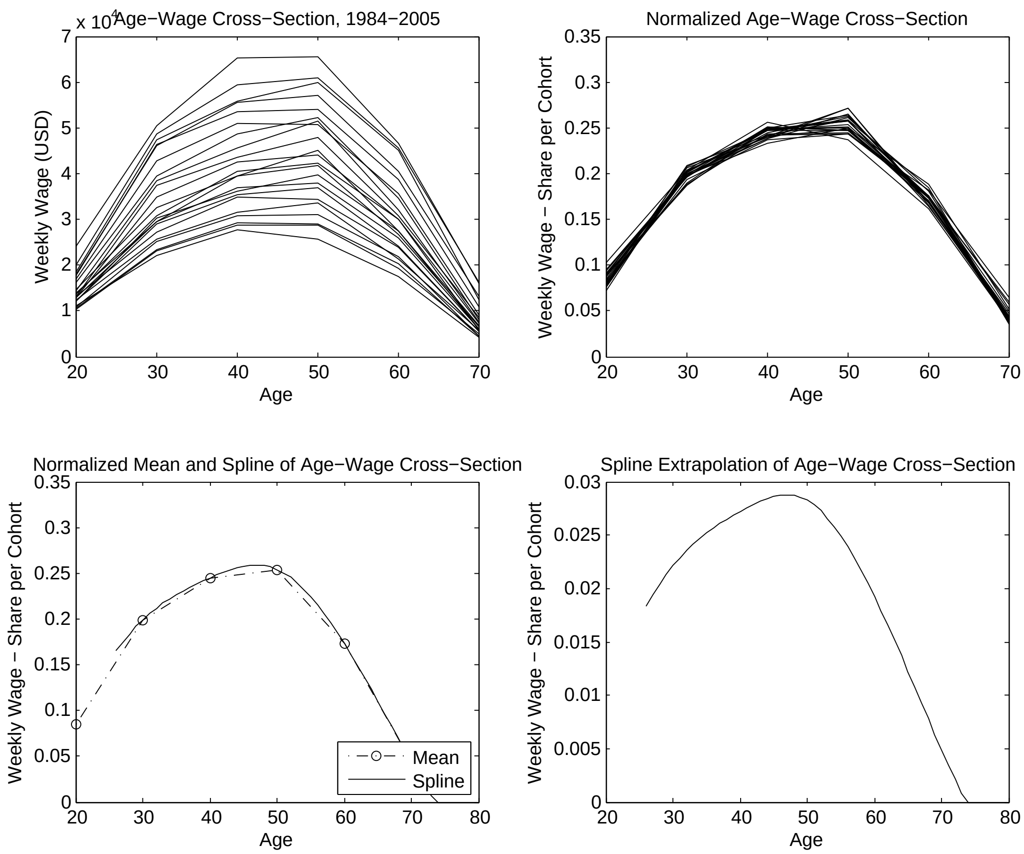

Figure 2.2: The household's effective labor units (l) over the life cycle

labor units over the life cycle is unaffected by growth. However, this may not be the case in reality, as the rapid technological changes associated with high growth rates may favor young workers, capable of quick learning and adjustment, over their older, more experienced counterparts. That being said, this aspect of the labor market is not explored in this paper.

Another caveat is the suspicion that surveys overstate average labor income for the last age bracket, since the number of observations for that category drops considerably, possibly causing a selection bias. It is likely that the drop in sample size is partly due to retirement, making the actual average labor income much lower, and thus causing an underestimation of household saving and the effect of growth. I choose not to try to adjust for that bias in order to avoid any suspicion that the results presented later are due to an over-correction on my part.

Additional parameters needing calibration are λ, κ and the international interest rate r.

λ, the probability of switching to an absorbing low growth regime, can have a range of values and may be country specific. A country close to the technology frontier is unlikely to sustain rapid growth for a very long time (and will thus have a higher λ). A country with strong, stable institutions is likely to remain in a high growth regime for a long time, i.e. have a lower λ. I use a several values of λ in my simulations: the benchmark case of 0%, and 2%, 5% and 10% cases. In addition, I set the per capita growth rate of the low growth regime to be 1%.

κ, the degree of credit constraint, can theoretically have any value in the $[0, -1]$ interval, with zero being a non-negative asset constraint, and -1 being no constraints at all. κ may be correlated to financial institution strength and sophistication, which may in turn be correlated with a country's level of development. I abstract away from such correlation, simply using several values of κ. The benchmark case is $\kappa = -1$, where no constraints are binding; I also use values $-0.1, -0.05, -0.02$ and 0.

r, the international interest rate, is assigned a value of 5%.

2.7 Simulation

2.7.1 The Life Cycle Baseline Model

The baseline model has balanced growth ($\gamma = 1$), no uncertainty ($\lambda = 0$) and no credit constraints ($\kappa = -1$). As figure 3 demonstrates, this simple life cycle model does generate a positive relation from growth to saving for the 0% to 4% per capita growth interval. However, slighly after 4% the relation becomes negative, with the saving rate at 9% roughly equal to that of 0%. This result is not in accordance with the data, which does not show the growth saving relation becoming negative (see Appendix II).

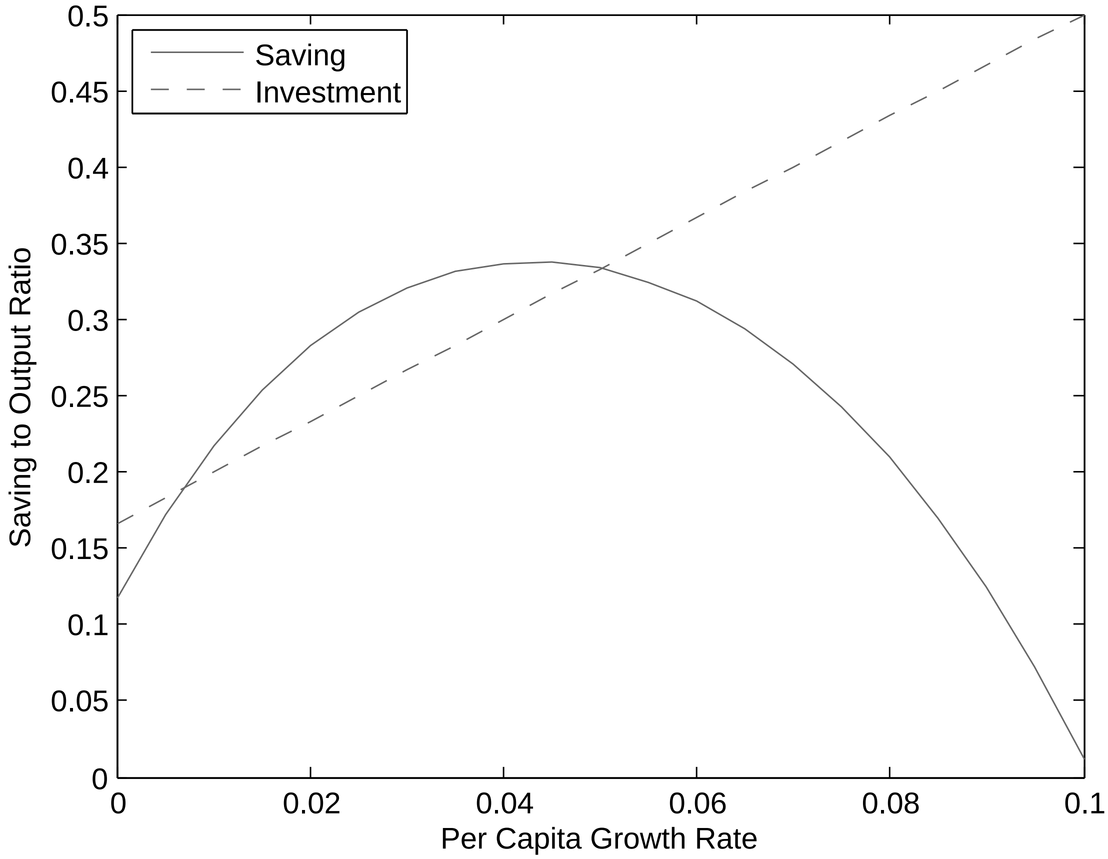

Figure 2.3: Saving and investment ratios as a function of the growth rate, baseline life cycle model

In conjunction with the investment line, which is based on equation (19), this saving curve predicts a positive current account for economies in the growth range of roughly 1% to 5% (for the given international interest rate and depreciation rate).

Breaking down the aggregate effect to the household level, we can see that growth causes the younger cohorts to dissave, much like the Carroll and Summers (1991) argument. As can be seen in figure 4, it also causes households to save more during middle-age years. Thus, in a high growth regime, consumption is relatively "'smoothed"' as households borrow substantially in their early, low income years, then save high amounts during middle-age years to pay back for early borrowing as well as amass assets for consumption during later years, when wage income dwindles down to nil.

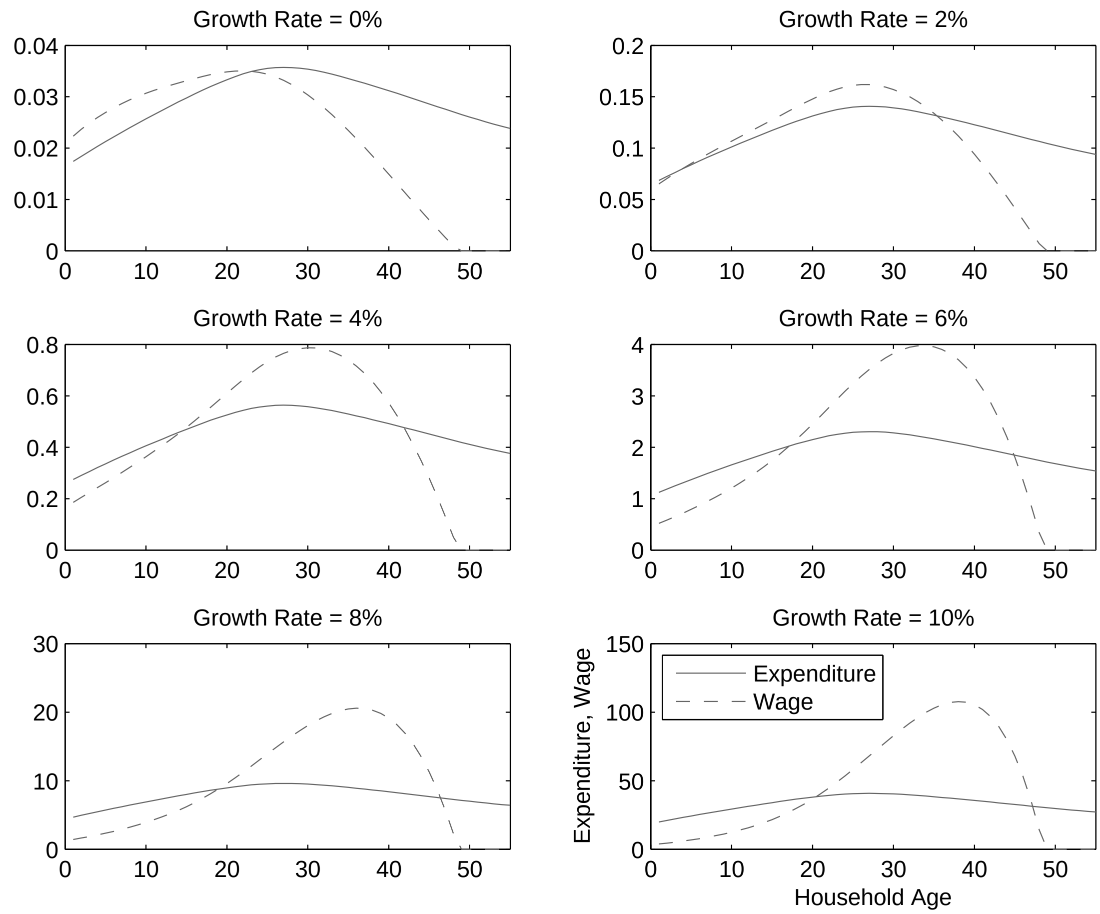

Figure 2.4: Household life cycle expenditure and wage, baseline model

Growth creates a wealth disparity between cohorts. Thus aggregate saving is influenced by the young generations dissaving much more than by the increased saving of the middle aged cohorts, as Figure 5 demonstrates, and can be seen to be grow up to the 4% growth regime and decrease thereafter.

In addition to failing to provide a persistent strong relation between growth and saving (on the aggregate level) over a range of high growth rates, the baseline model also does not predict an increasing path of household consumption over the lifetime in high growth regimes. This correlation between higher household consumption expenditure growth rates correlated and higher aggregate growth is demonstrated in both Carroll and Summers (1991) and Lee et al. (2006).

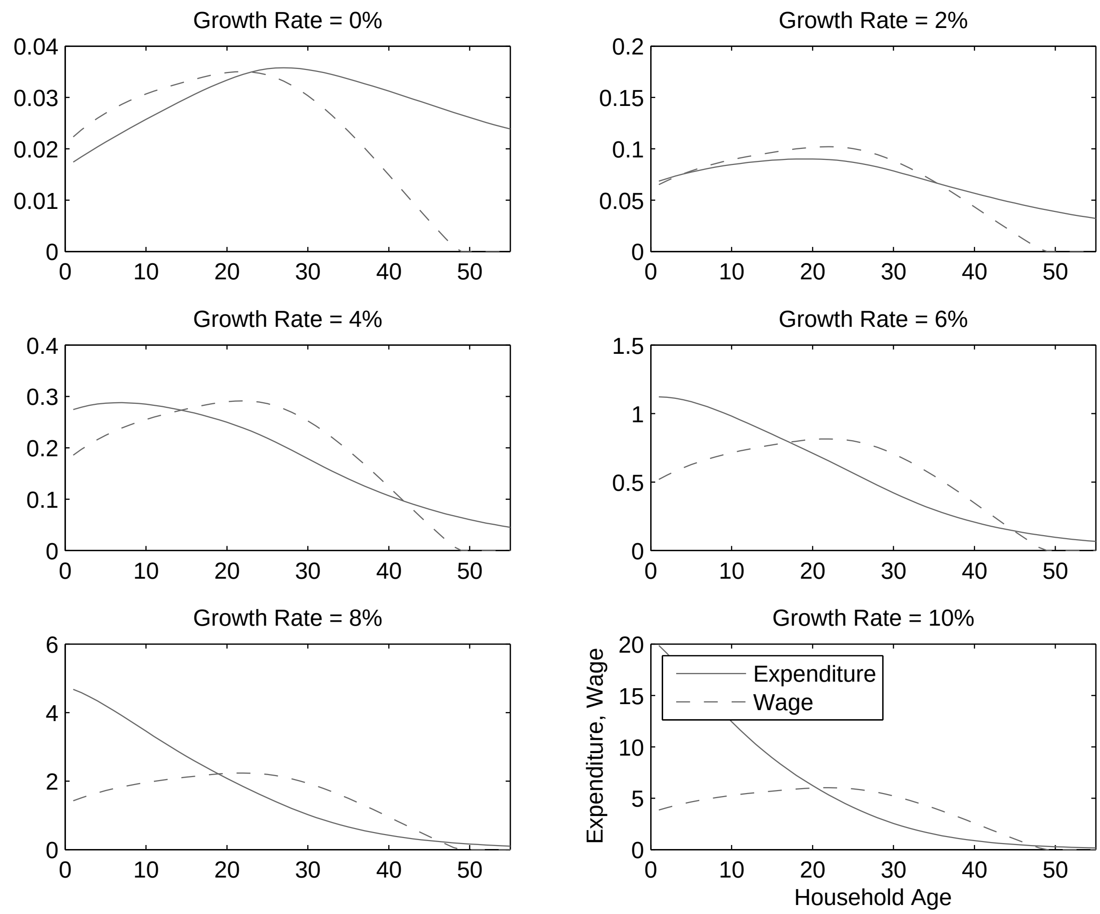

Figure 2.5: Cross section of household expenditure and wage, baseline model

Another counter-factual result of the simulation is the massive gap between the consumption of young and old cohorts. For example, under a regime of 8% per-capita growth, a household of cohort 1 (the head of the household is 25-26) has over 10 times the consumption of a household of cohort 40 (who's head is 59-60). Empirically, as Carroll and Summers (1991) and Lee et al. (2006) show, no such gap exists, and the old cohorts of a rapidly growing economy are no worse off compared to younger cohorts than those of a slowly growing economy.

2.7.2 Structural Change

With γ values of less than one we can simulate the effects of a price level increase associated with unbalanced growth. The simulation is performed for an economy with a current traditional sector share of one half.

The ensuing increasing price level does indeed lead the households to consume a larger proportion of their income during later years, dissaving less when young and saving more in middle-age years, as can be seen in Figure 6.

This effect is strong enough to substantially change the aggregate picture, as Figures 7 and 8 demonstrate. Comparing Figure 7 to Figure 5, we can see that the disparity in consumption between young and old cohorts for the high growth regimes is somewhat reduced.

Figure 8 shows the negative effect of growth on the aggregate saving rate is very much reduced by introducing structural change. The effect remains positive past the 6% growth mark, and the decreasing trend for higher growth is much flatter, resulting in saving rates for high growth values being markedly higher than for the low values.

Structural change allows the simulation to grow closer to some of the stylized facts: high growth economies enjoy high rates of saving, the expenditure gap between young and old cohorts (1 and 40) has diminished greatly, and life cycle expenditure profiles do seem to be increasing in age for high growth values.

However, while it does seem to move in the right direction, structural change in itself does not seem sufficiently strong to completely match the stylized facts: as Figures 18-20 in Appendix II show, the data does not reflect a decrease in the saving rate for high growth rates. In addition, the simulated expenditure cross sections still exhibit too much disparity, in contradiction of the empirical evidence (e.g. Carroll and Summers (1991)).

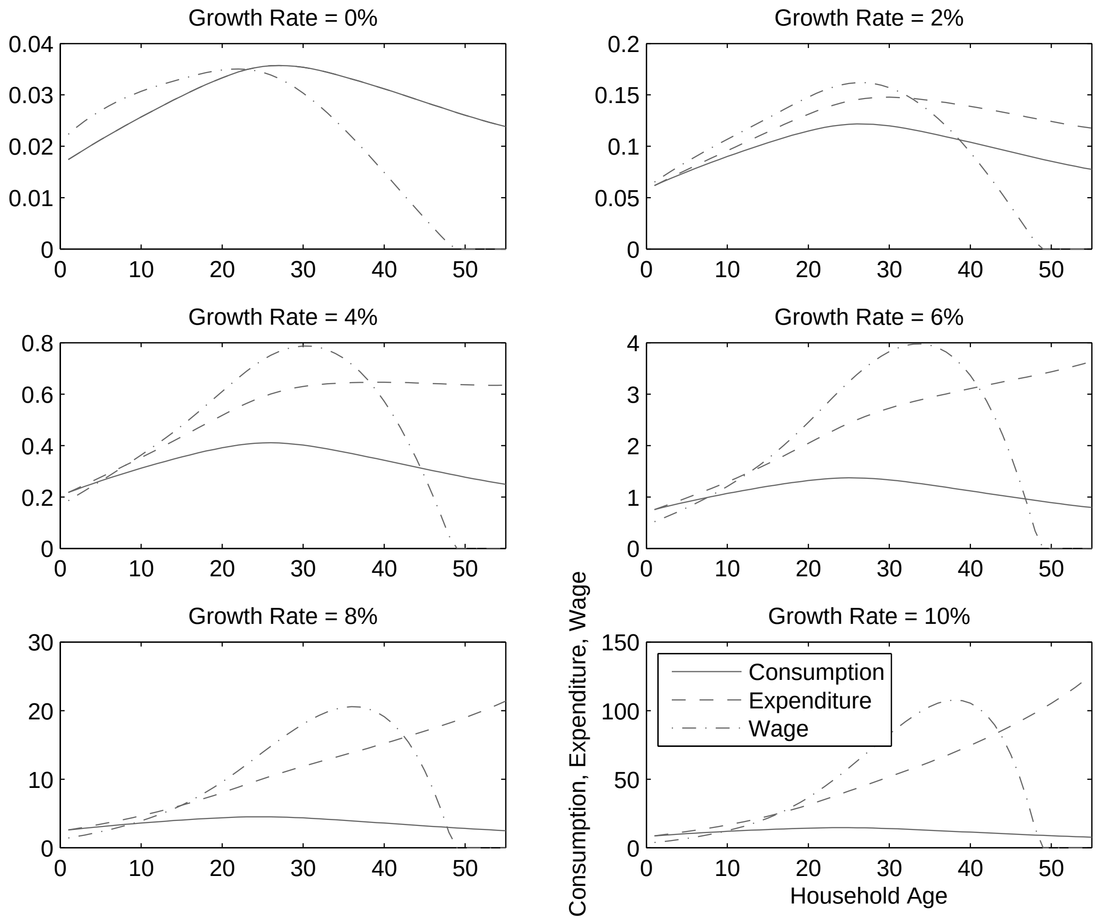

Figure 2.6: Household life cycle consumption, expenditure and wage, with structural change ($\gamma = 0.5$)

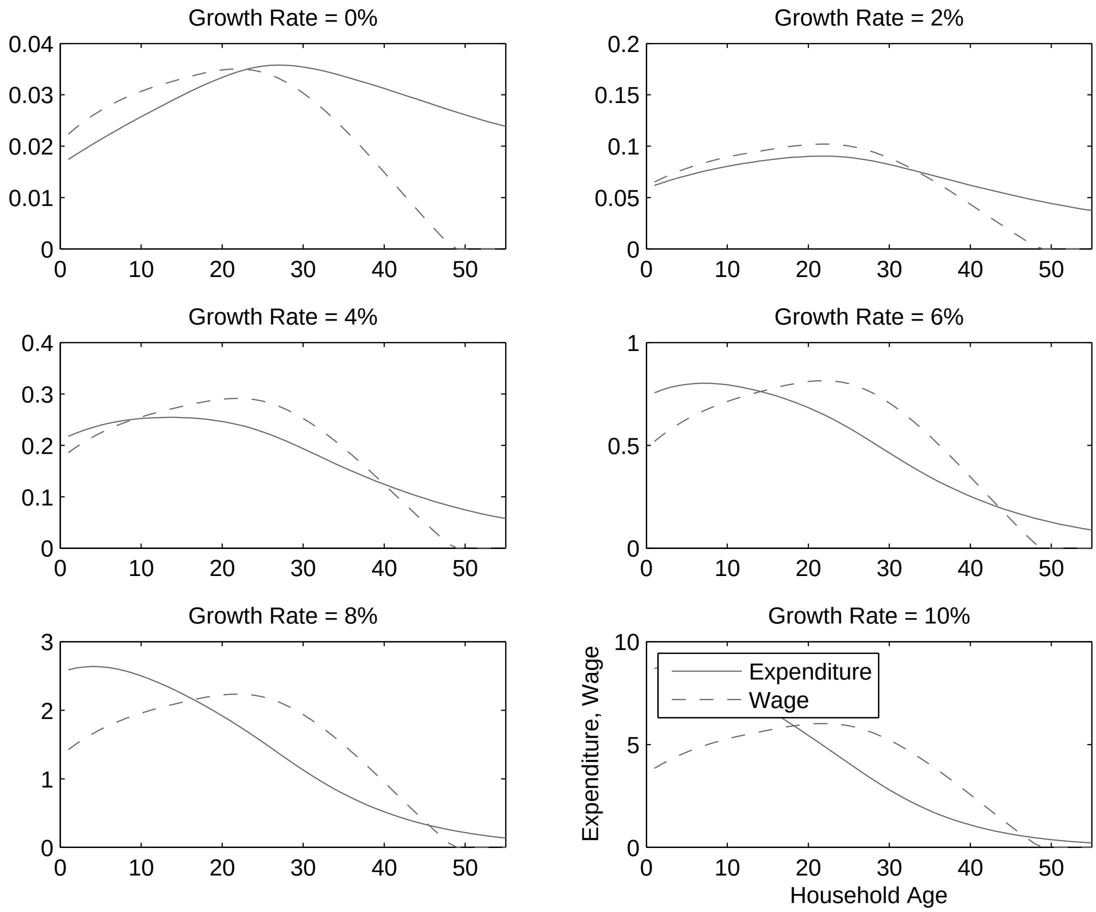

Figure 2.7: Cross section of household expenditure and wage, with structural change ($\gamma = 0.5$)

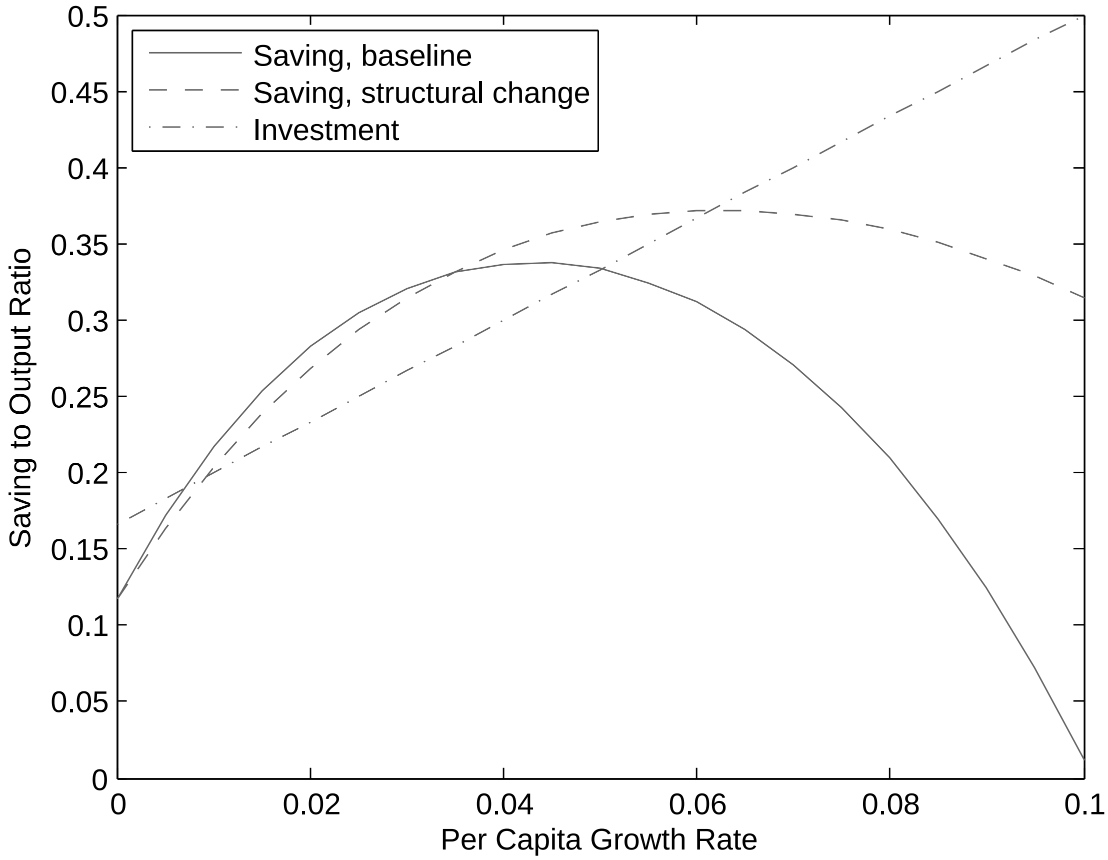

Figure 2.8: Saving and growth with structural change ($\gamma = 0.5$)

2.7.3 Growth Uncertainty

Under growth uncertainty there is a constant probability λ that the economy will drop from its high growth state to an absorbing low growth state. The effect of this drop is of course stronger the larger is the gap between the two states, e.g. a drop from 8% annual per capita growth to 0% is obviously to be dreaded much more than a drop from 4% to 2%. In this simulation the absorbing low growth regime is assumed to be 1% per capita output growth.

Since the state of interest here is that of prolonged high growth, then regardless of the size of λ, the drop to low growth never actually occurs. In other words, while households are forced to plan for low growth eventualities, we are interested only in observing the high growth realizations of this stochastic process. The effect on household saving decisions is twofold: on the one hand it is forced to save toward a future where its income is no longer growing like it used to. This is reminiscent of a precautionary saving motive. On the other hand, since we are observing only the high growth realizations, the household's expected future income is constantly updated upwards, resulting in an ex-post sub-optimal expenditure path where it is borrowing less than it would had its high future income been certain.

As Figure 9 shows, a λ of 5% greatly reduces dissaving during early years in the model. However, the magnifying effect of growth (i.e. young household being so much wealthier) is strong enougn that this dissaving eventually overwhelms the increased saving of middle-aged cohorts, as Figure 10 shows, though at higher growth rates than the baseline case, as shown in Figure 11.

Notice also that comparing Figures 4 and 9, it is evident that at 2% growth uncertainty only generates an extermely minute difference in the household expenditure path. This goes to show that not including the uncertainty element in the process of calibrating q should produce no bias.

It can be seen that the presence of uncertainty can be very influential in increasing household saving and decreasing the tendency of the younger cohorts to dissave. A risk of $\lambda = 2\%$, which is extremely low considering it results in an average high growth period of 50 years, is enough to greatly reduce the negative effect of growth on saving in high growth rates. Higher levels of risk further flatten the growth-saving graph (Figure 11), so that, while still decreasing after some point, the simulated saving rate remains much higher for the high values of growth than for the low values.

In addition, growth uncertainty further reduces the expenditure gap between young and old cohorts. For example, the ratio of the expenditure of cohorts 1 and 40 is reduced to to 3 to 1 under 8% growth. It does not, however, provide increasing life

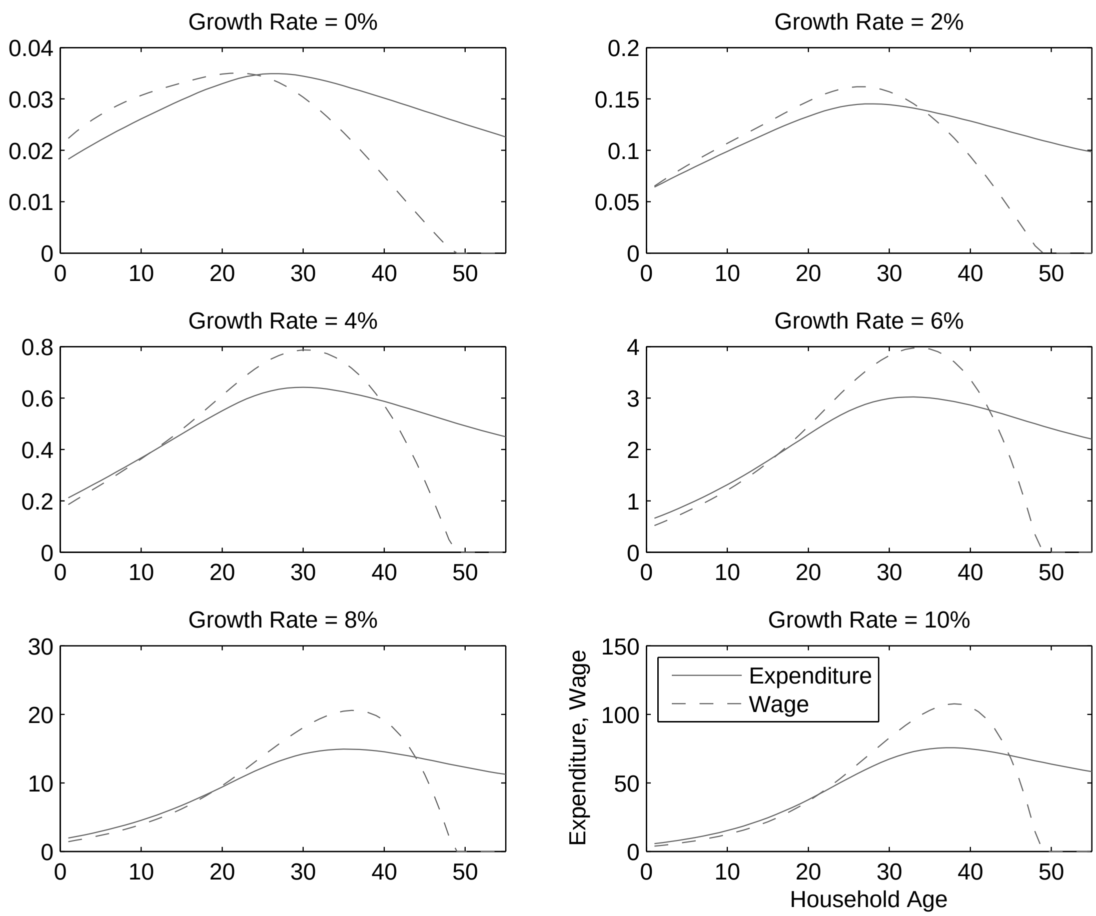

Figure 2.9: Household life cycle expenditure and wage, with growth uncertainty ($\lambda = 5\%$)

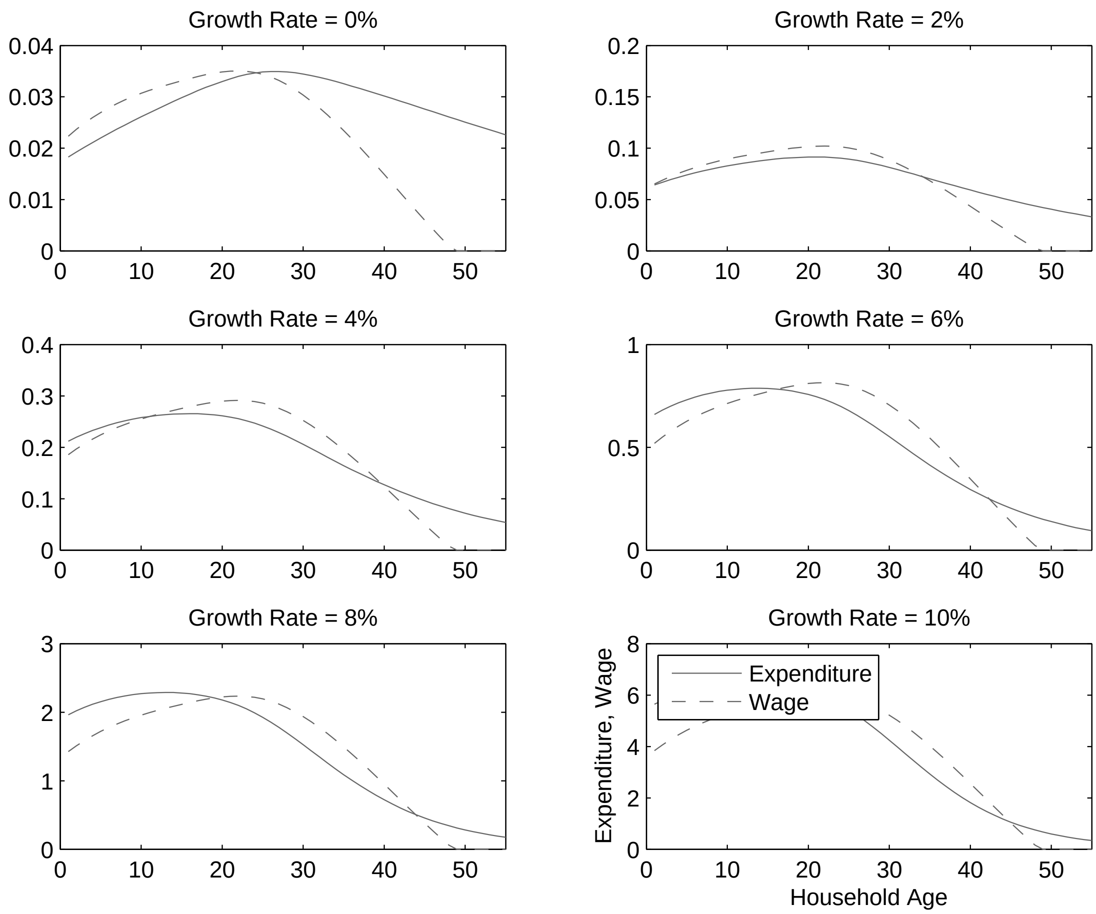

Figure 2.10: Cross section of household expenditure and wage, with growth uncertainty ($\lambda = 5\%$)

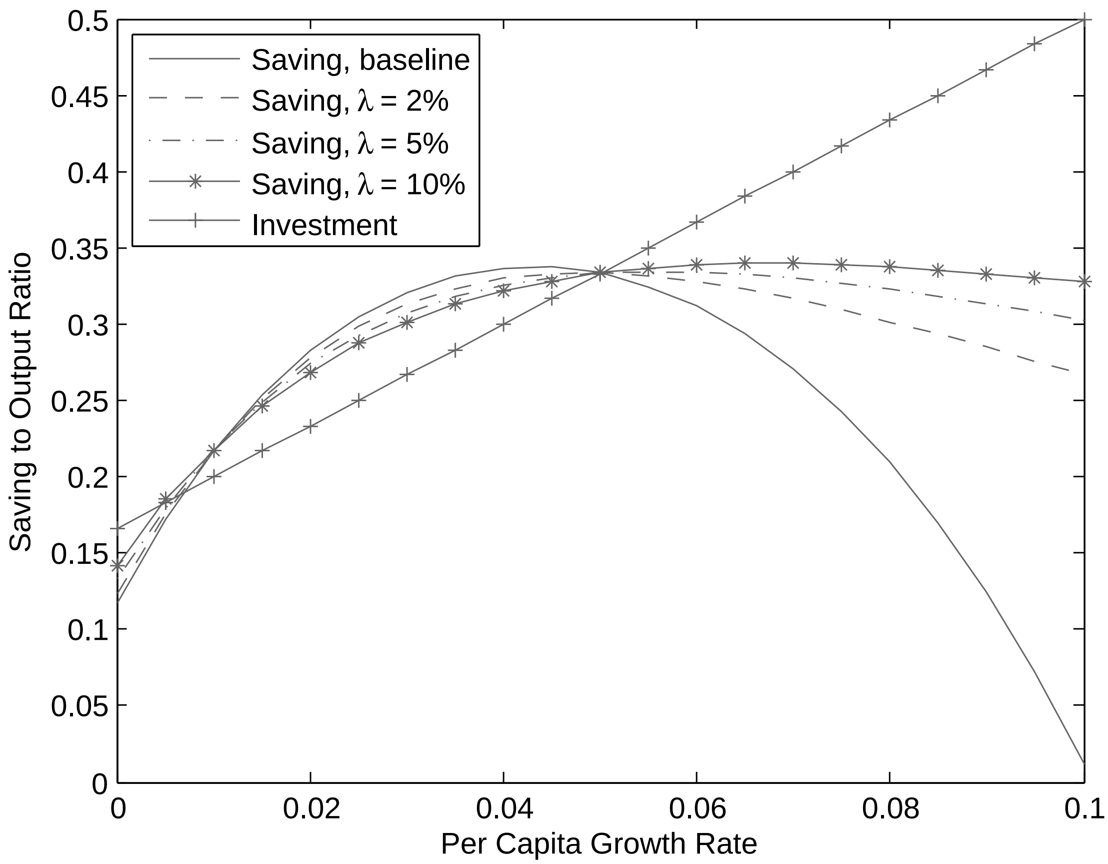

Figure 2.11: Saving and growth - a sensitivity analysis for different levels of growth uncertainty

cycle expenditure profiles.

2.7.4 Credit Constraints

Credit constraints, inhibiting the ability of younger generations to dissave, can have a powerful impact on aggregate saving. As Figure 12 shows, the constraint needs to be quite tight for there to be a significant effect, as $\kappa = -0.1$ is almost identical to $\kappa = -1$ (the no constraint case). Clearly once the younger cohorts are able to borrow 10% of their future wealth the constraints are no longer binding. A non-negativity constraint ($\kappa = 0$) is shown in Figure 12 to create a positive effect of growth on saving all the way up to 10% per capita growth (the upper limit of the simulation).

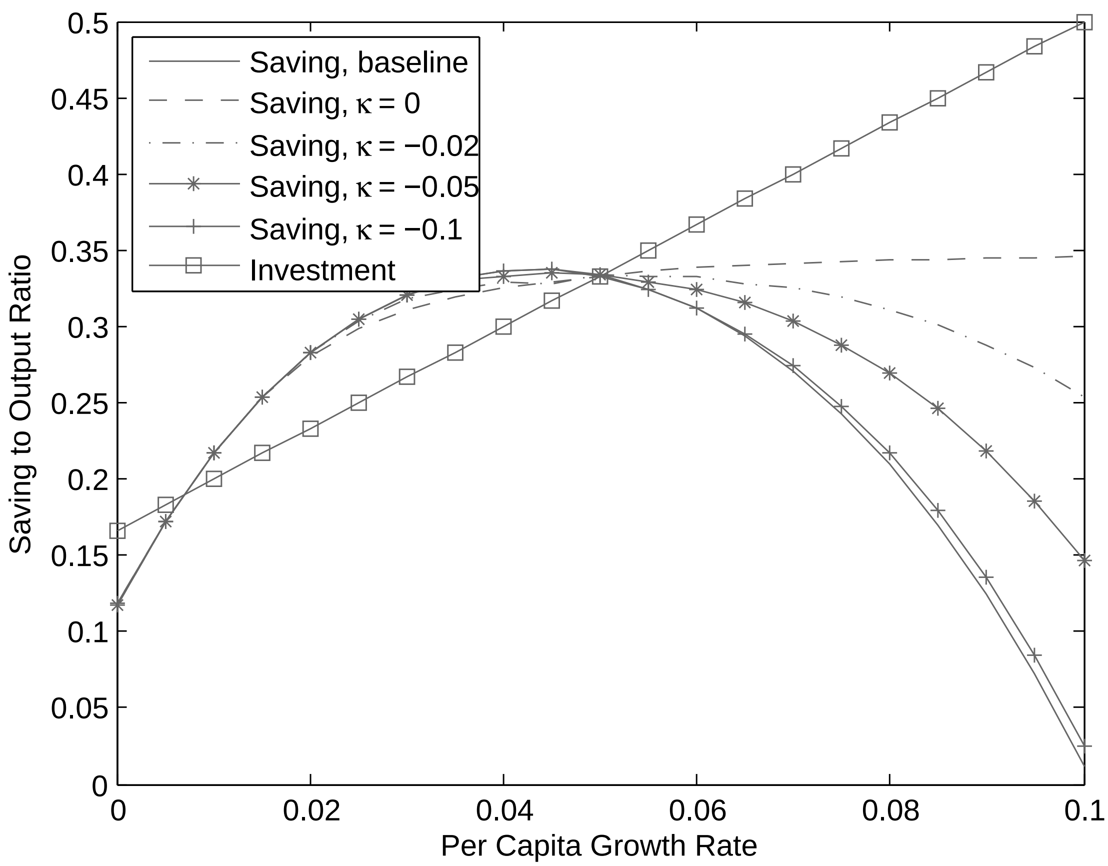

Figure 2.12: Saving and growth - a sensitivity analysis for different levels of credit constraints

Figure 13 shows the effects of a non-negativity constraint ($\kappa = 0$) on the household's life cycle expenditure. The higher is growth, the longer is the household bound by the constraint. With younger cohorts unable to borrow, the amplifying effect of growth makes the savings of middle-aged cohorts outweigh the dissavings of older cohorts. Notice however that credit constraints element does not lead to an increasing path of life cycle consumption, and thus falls short of matching the data on its own.

Notice also that at 2% per capita growth the range in which the constraints hold is very minimal. Since they do not seem to affect the expenditure path at this growth rate, it is superfluous to include credit constraints in the q calibration process.

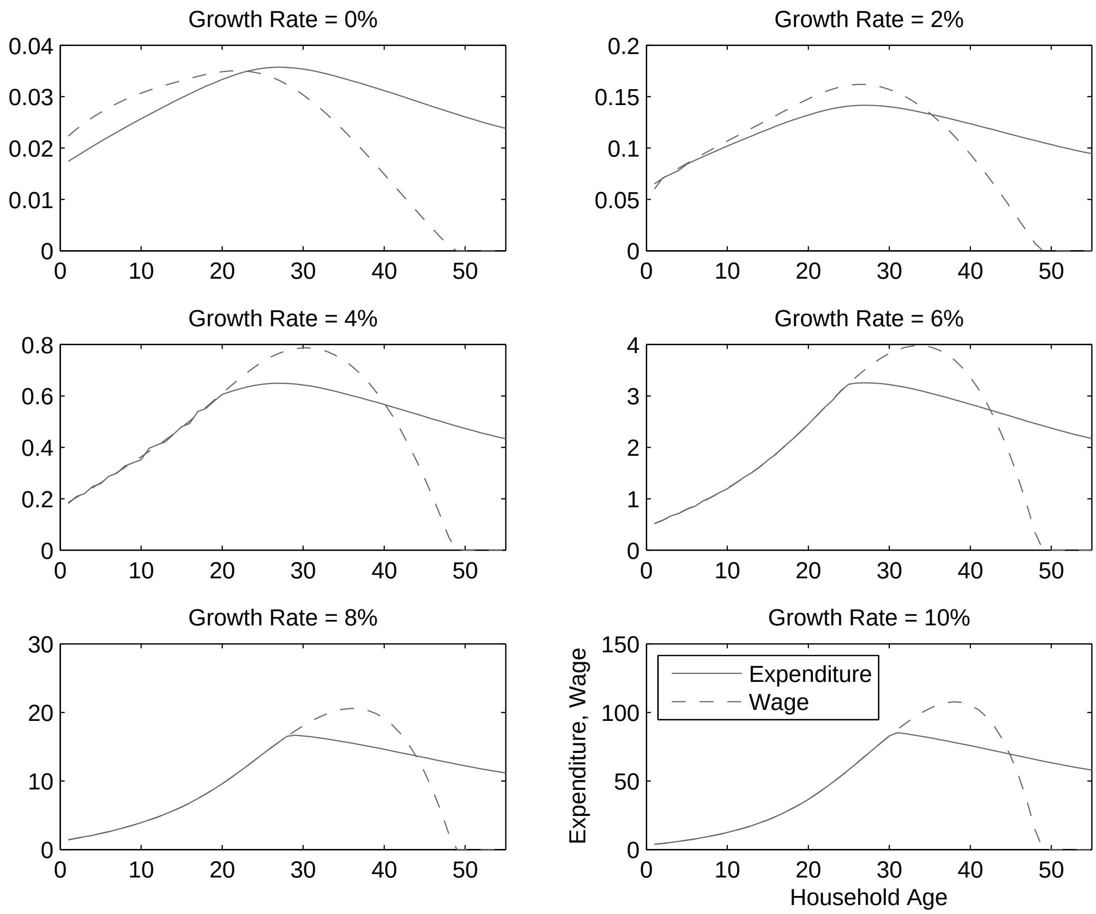

Figure 2.13: Household life cycle expenditure and wage, with credit constraints ($\kappa = 0\%$)

As Figure 14 also shows, credit constraints do decrease the simulated intra cohort consumption disparity. For example, the ratio of period 1 cohorts and period 40 cohorts expenditure under a growth rate of 8% to be about 2 to 1.

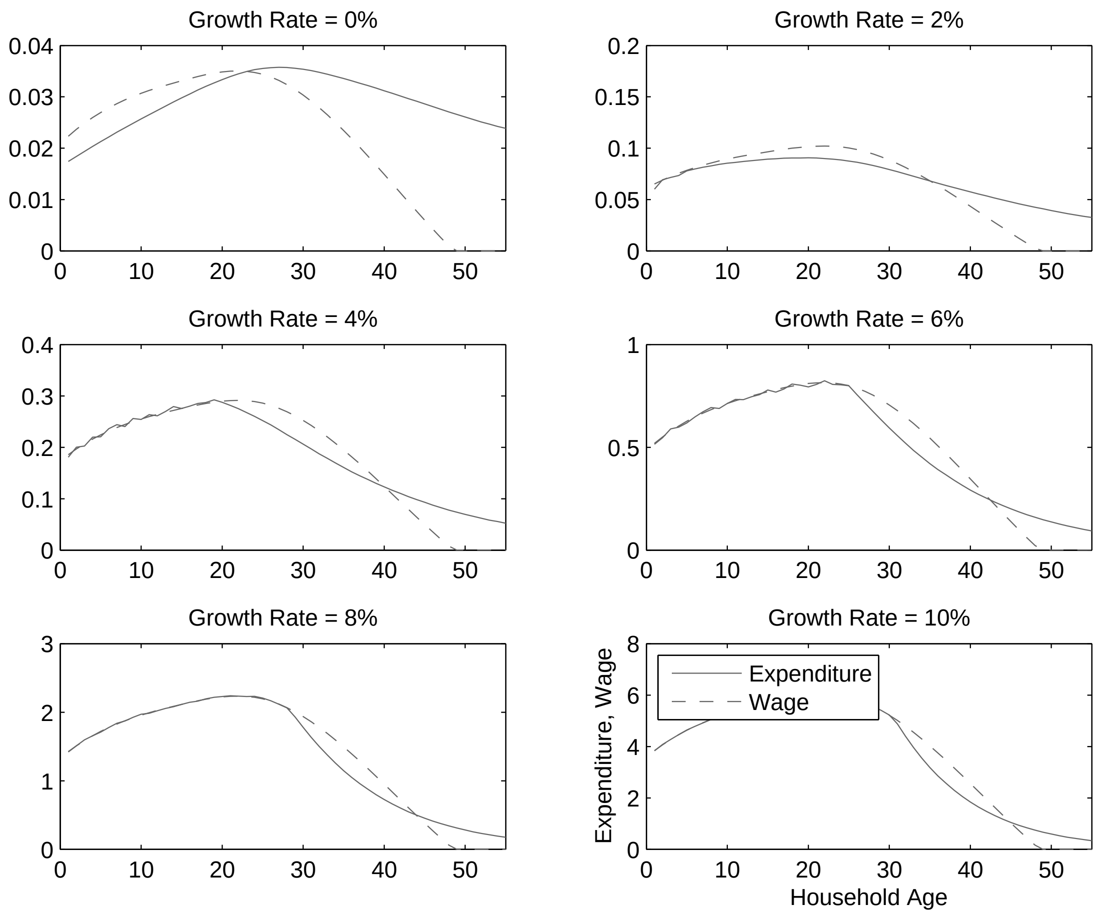

Figure 2.14: Cross section of household expenditure and wage, with credit constraints ($\kappa = 0\%$)

2.7.5 Combining the Three: Structural Change, Growth Uncertainty and Credit Constraints

Figure 15 presents a comparison of the effects of various combinations of the three elements used in this paper in conjunction with the life cycle model - structural change, growth uncertainty and credit constraints. Clearly, including non-negativity credit constraints ($\kappa = 0$) with any other elements results in a positive effect of growth on the saving rate in all of the simulated range (0% to 10% per capita output growth). Interestingly, using structural change and growth uncertainty without credit constraints ($\gamma = 0.5, \lambda = 5\%$ and $\kappa = -1$) creates a positive growth-saving correlation for most of the range with a very mild negative relation towards the upper end. This can hardly be said to be in constradiction with the data.

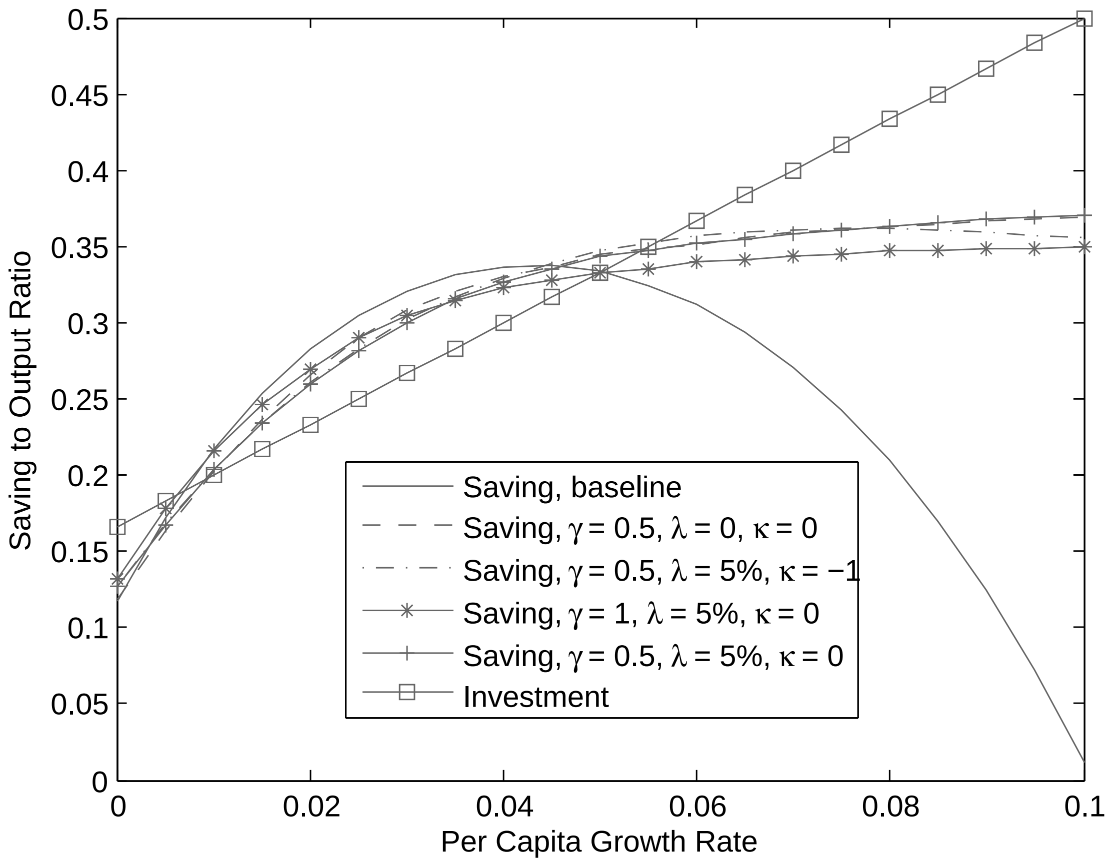

Figure 2.15: Saving and growth - simulated combinations of structural change, growth uncertainty and credit constraints

However, while many of the possible combinations produce a reasonable growth-saving correlation, only the combinations including the structural change element are able to mimic the stylized fact of a correlation between household expenditure growth and aggregate growth. Figure 16 serves as an example.

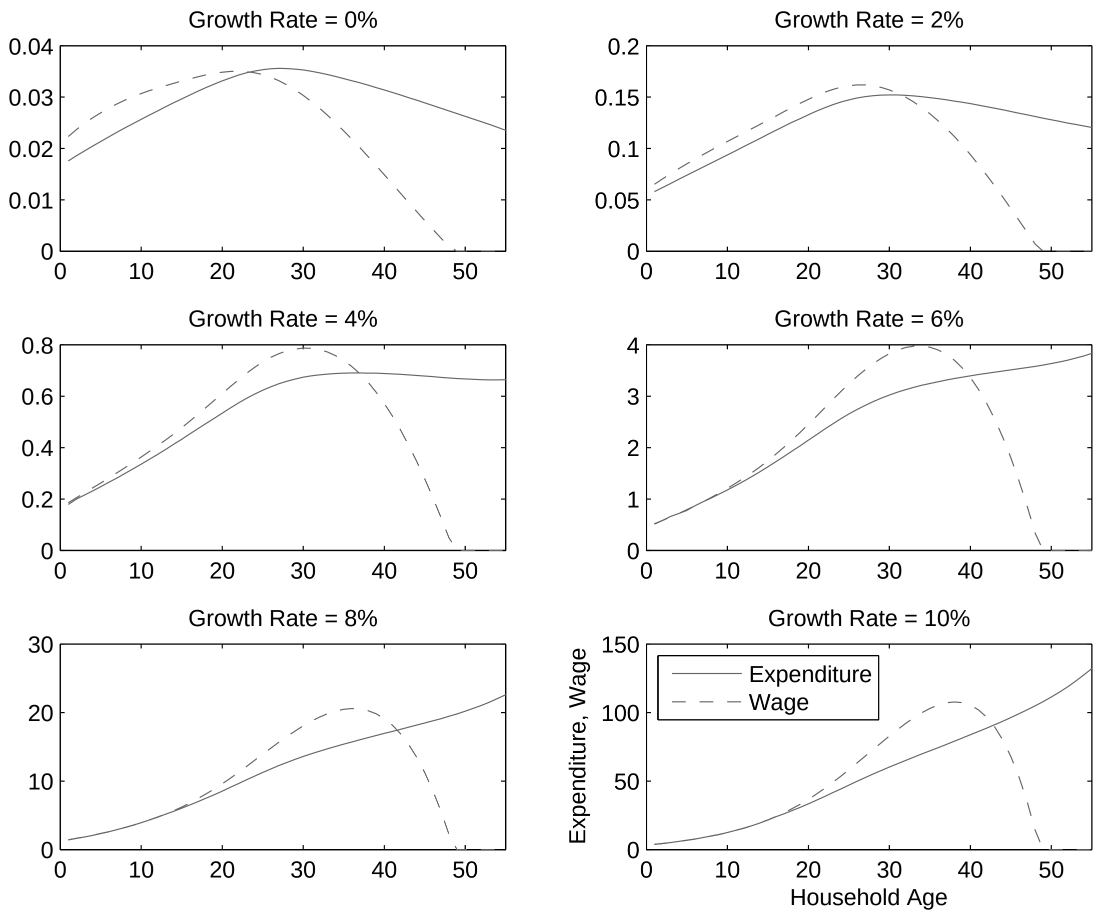

Figure 2.16: Household life cycle expenditure and wage - cumulative effect of structural change, growth uncertainty and credit constraints

Also of note is the fact that combining the three elements further reduces the cross cohort expenditure disparity (see Figure 17).

The simulated cross sectional disparity is still greater the higher is growth, a fact not in line with the data (e.g. Carroll and Summers (1991)). While this is not within the scope of this paper, I believe that intergenerational transfers could play a role in "flattening" that disparity.

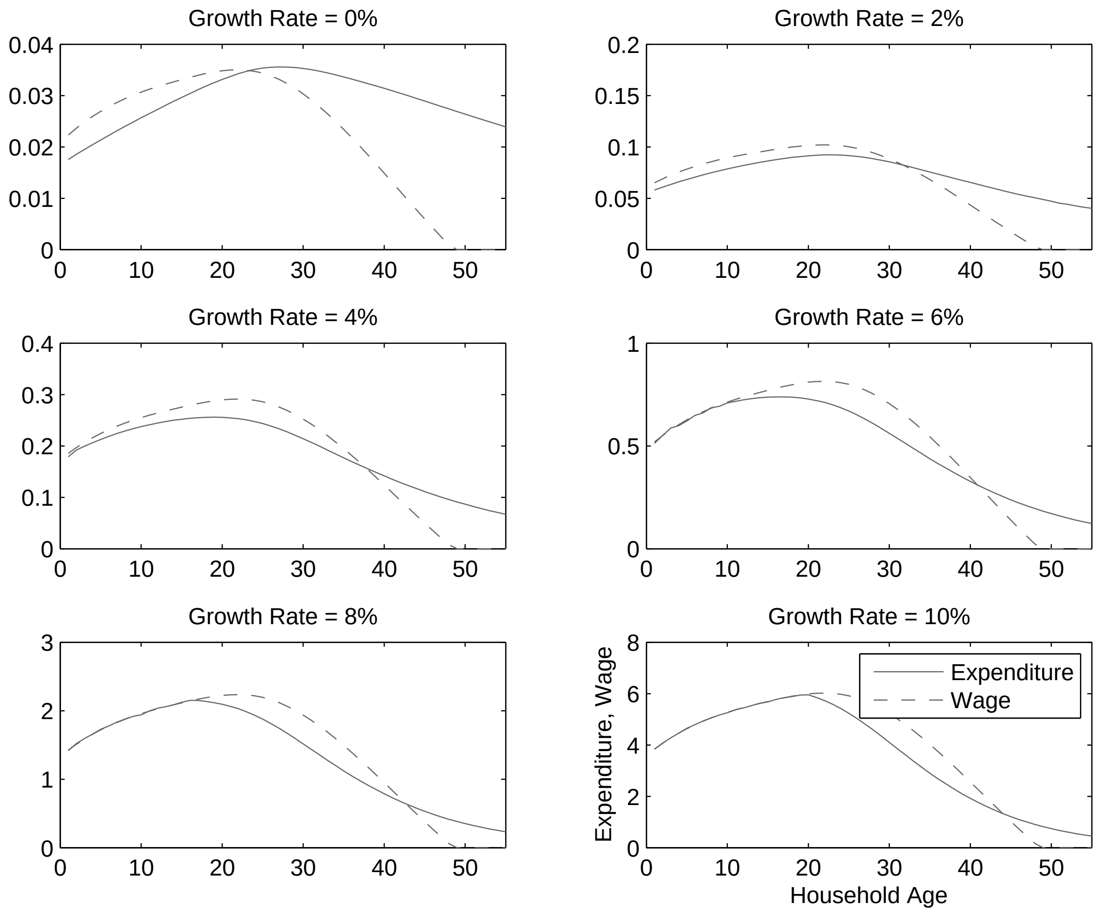

Figure 2.17: Cross section of household expenditure and wage - cumulative effect of structural change, growth uncertainty and credit constraints

2.8 Conclusion

The addition of structural change to a Modigliani-like life-cycle model does indeed strengthen its ability to simulate the widely observed phenomenon of high saving rates in rapidly growing countries. When adding the elements of growth uncertainty and credit constraints (together or separately), the model is able to create positive causality from growth to saving, while allowing growth to fully affect the households income path, overcoming the critique made in Carrol and Summers (1991). In addition, structural change presents itself as a viable explanation for the stylized fact of correlation between aggregate growth and the growth in consumption expenditure during the household life cycle, and in conjunction with the other elements greatly reduces the simulated cross cohort expenditure disparity resulting from rapid growth.

Growth is driven by improving productivity in the modern sector, changing the relative price of the goods produced in the traditional sector. In essence, realizing that future consumption of traditional goods (e.g. housing, health, education etc.) is getting expensive, households choose to save more for their later years, when they cannot hope for much labor income despite continued aggregate growth. This serves to increase the aggregate saving rate for a range of growth values, for which it can be said that the life cycle substitution effect overpowers the life cycle income effect.

In the international context, explaining the high saving rates observed in rapidly growing developing countries also serves toward getting to the root of the direction of capital flows in the world's economy. The excess saving created by growth allows poor but rapidly developing countries to maintain a trade surplus and even lend to richer but slower growing economies. With large parts of the world (China, India etc.) poised for a prolonged catching up process, this trend may dominate international capital markets over the next several decades.

2.A Algebraic Derivations

Utility maximization for a household created at period t with a given an output growth rate of g:

$$L = \sum_{j=1}^{J} \beta^{j-1} q(j) \frac{c(j,t+j-1,g,r)^{1-\theta}-1}{1-\theta} +$$
$$\lambda(t,g,r) \sum_{j=1}^{J} (1+r)^{-j+1} \left[w(j,t+j-1,g) - P(t+j-1,g)\, c(j,t+j-1,g,r) \right]$$

First Order Conditions:

$$c(j,t+j-1,g,r) : \beta^{j-1} q(j)\, c(j,t+j-1,g,r)^{-\theta} - \lambda(t,g,r)(1+r)^{-j+1} P(t+j-1,g) = 0$$
$$\lambda(t,g,r) : \sum_{j=1}^{J} (1+r)^{-j+1} \left[w(j,t+j-1,g) - P(t+j-1,g)\, c(j,t+j-1,g,r) \right] = 0 \tag{2.34}$$

Calculating the household's consumption path:

$$c(j,t+j-1,g,r) :$$
$$\beta^{j-1} q(j)\, c(j,t+j-1,g,r)^{-\theta} - \lambda(t,g,r)(1+r)^{-j+1} P(t+j-1,g) = 0$$
$$\Rightarrow c(j,t+j-1,g,r) = \left\{ \frac{q(j)}{\lambda(t,g,r) P(t+j-1,g)} [\beta(1+r)]^{j-1} \right\}^{\frac{1}{\theta}}$$
$$c_0(t,g,r) \equiv \lambda(t,g)^{-\frac{1}{\theta}} \Rightarrow c(j,t+j-1,g,r) = c_0(t,g,r) \left\{ \frac{q(j)}{P(t+j-1,g)} [\beta(1+r)]^{j-1} \right\}^{\frac{1}{\theta}}$$
$$\lambda(t,g,r) :$$
$$\sum_{j=1}^{J} (1+r)^{-j+1} w(j,t+j-1,g) = \sum_{j=1}^{J} (1+r)^{-j+1} P(t+j-1,g)\, c(j,t+j-1,g,r)$$
$$= c_0(t,g,r) \sum_{j=1}^{J} (1+r)^{-j+1} P(t+j-1,g)^{\frac{\theta-1}{\theta}} \left\{ q(j) [\beta(1+r)]^{j-1} \right\}^{\frac{1}{\theta}}$$
$$\Rightarrow c_0(t,g,r) = \frac{\sum_{j=1}^{J} (1+r)^{-j+1} w(j,t+j-1,g)}{\sum_{j=1}^{J} (1+r)^{-j+1} P(t+j-1,g)^{\frac{\theta-1}{\theta}} \left\{ q(j)[\beta(1+r)]^{j-1} \right\}^{\frac{1}{\theta}}} \tag{2.35}$$

Substituting for w:

$$\begin{aligned} c_0\left(t,g,r\right) &= (1+g)^t \frac{\sum\limits_{j=1}^{J}\left(\frac{1+g}{1+r}\right)^{j-1} l(j)W(j-1,g)}{\sum\limits_{j=1}^{J}(1+r)^{-j+1}P(t+j-1,g)^{\frac{\theta-1}{\theta}}\left\{q(j)[\beta(1+r)]^{j-1}\right\}^{\frac{1}{\theta}}} \\ &= (1+g)^t \frac{\Omega(0,g,r)}{\sum\limits_{j=1}^{J}(1+r)^{-j+1}P(t+j-1,g)^{\frac{\theta-1}{\theta}}\left\{q(j)[\beta(1+r)]^{j-1}\right\}^{\frac{1}{\theta}}} \\ &\Rightarrow c\left(j,t+j-1,g,r\right) = (1+g)^t \frac{\Omega(0,g,r)\left\{\frac{q(j)}{P(t+j-1,g)}[\beta(1+r)]^{j-1}\right\}^{\frac{1}{\theta}}}{\sum\limits_{j=1}^{J}(1+r)^{-j+1}P(t+j-1,g)^{\frac{\theta-1}{\theta}}\left\{q(j)[\beta(1+r)]^{j-1}\right\}^{\frac{1}{\theta}}} \end{aligned} \tag{2.36}$$

Aggregate saving:

$$S\left(t,g,r\right) = Y\left(t,g\right) - P\left(t,g\right)\sum_{j=1}^{J} c\left(j,t,g,r\right) \tag{2.37}$$

Saving to output ratio:

$$s(g,r) \equiv \frac{S}{Y}\left(g,r\right) = 1 - \frac{P\left(t,g\right)\sum\limits_{j=1}^{J} c\left(j,t,g,r\right)}{Y\left(t,g\right)} \tag{2.38}$$

2.B Growth and Saving Scatter Plots

Data is from WDI. On the X axis: average per capita growth for years 1976-1985, 1986-1995 and 1996-2005. On the Y axis: gross saving rate for years 1985, 1995 and 2005. Growth rates are averaged out over a decade since they are much more volatile than saving rates. I am using the saving rate at the end of the decade because growth seems to affect saving with a lag of a few years.

Country selection – all countries meeting the following criteria:

At least 9 data periods in the given time period.

Both saving and domestic saving data available.

Variance of growth in the given time period is less than 25. This is in order to get rid of noisy observations of countries which are very volatile (off equilibrium for some reason).

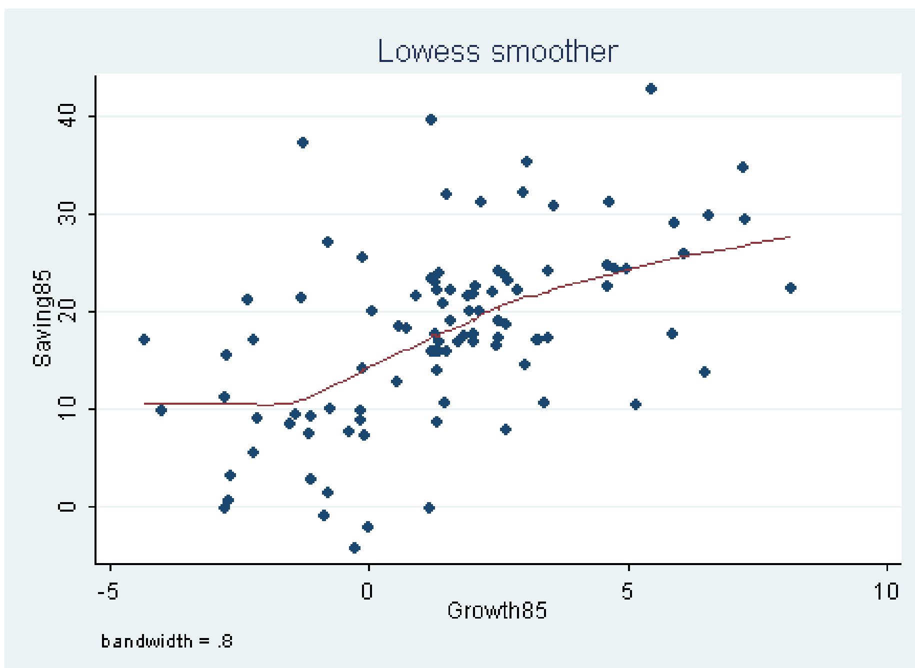

Figure 2.18: Saving (1985) vs. Growth (1976-1985 average)

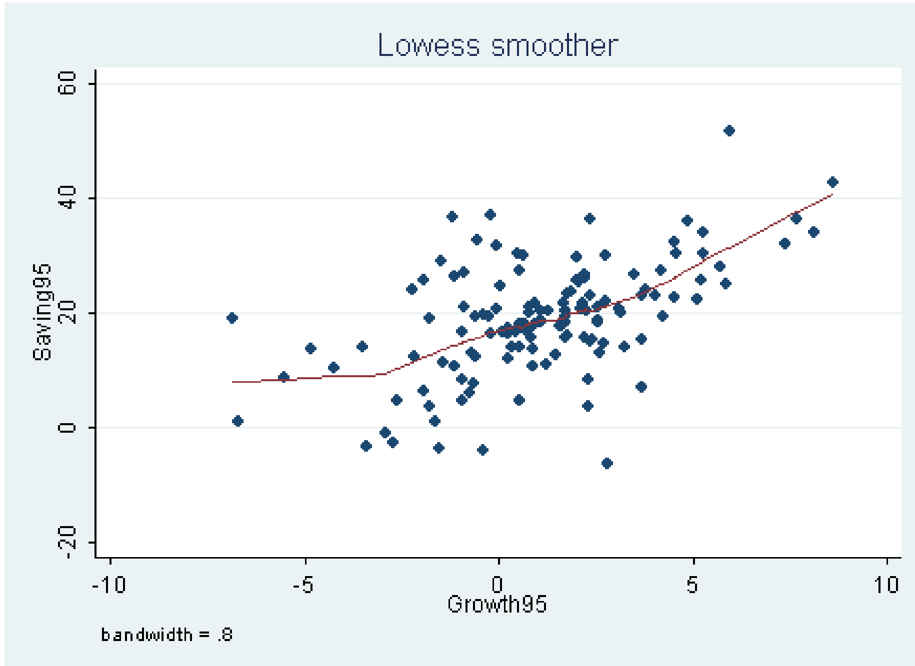

Figure 2.19: Saving (1995) vs. Growth (1986-1995 average)

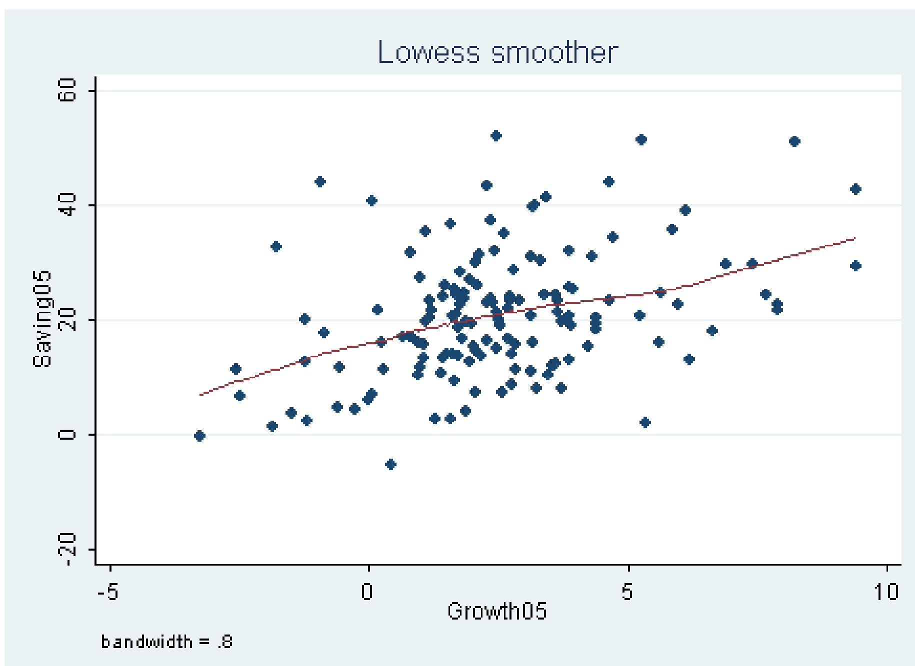

Figure 2.20: Saving (2005) vs. Growth (1996-2005 average)

CHAPTER III

Future shock: the effect of rapid growth on earning ability with age

"To assert that man must adapt seems superfluous. He has already shown himself to be among the most adaptable of life forms. He has survived Equatorial summers and Antarctic winters. He has survived Dachau and Vorkuta. He has walked the lunar surface. Such accomplishments give rise to the glib notion that his adaptive capabilities are 'infinite.' Yet nothing could be further from the truth. For despite all his heroism and stamina, man remains a biological organism, a 'biosystem,' and all such systems operate within inexorable limits."

– Alvin Toffler, Future Shock (1970)

3.1 Introduction

In his book, Alvin Toffler describes a future in which the pace of technological change is rapid enough to strain individuals and social structures. To a certain extent, this is something that we see all around us: as technology progresses, the young are more able to pick up new skills and adapt, and the technological divide between generations grows.

In this paper I examine the effect of an even greater rate of change: that of the 'catching-up' process, in which select developing countries catch up with the developed world, on the earning ability of older cohorts. The 'catching-up' process is characterized by rapid economic growth which in turn is fueled by a massive technology adoption, reshaping of industries and changing social structures. Does this process also create a divide between generations? Are the old less able to adapt new technologies and take advantage of new opportunities? More concisely, is the relative earning ability of older cohorts compared to younger ones affected by rapid growth?

This question, closely resembling that of Laitner and Stolyarov (2005), has yet to receive a definite answer.

Investigating the effect of prolonged rapid growth, i.e. a 'catching-up' process, on the earning-age profile, I show, using Chinese urban household survey data, that the earning-age profile of individuals in high growth locations to be distinctly different from those in low growth ones. High growth seems to decrease the earning ability of older cohorts compared to the young, resulting in a decreasing earning-age cross-sectional profile.

3.2 Analyzing 'Future Shock'

Chamon and Prasad (2010) and Song and Yang (2010) show Chinese income-age cross-sections to dramatically change over time: from a fairly standard profile in the early 1990's, where income increases with age until the mid 50's and then drops sharply toward retirement, to a much flatter profile late in the 2000's. Chamon and Prasad (2010) actually shows a partially decreasing profile that peaks around the late 20's.

I argue that this change is due to the effect of the rapid growth 'future shock' in China, in conjunction with the increasing share of private employment. But through what channels is this shock affecting the earning ability of older-age workers? The immediate suspects can be summed up as following:

- Obsolete education – new times require new skills, and rapid growth is likely to change both the preferred fields of study and the contents learned within each field. Thus possession of current education gives the young an advantage in the workplace. [1]

- Depreciating experience – knowledge acquired on the job is less valuable as the job itself changes, making the difference between 5 years and 25 years of experience less significant.

- Slower learning – the young are generally quicker and more willing to adapt to new circumstances, a great advantage in a rapidly changing environment.

In short, if we accept that a person's wage is highly dependent on human capital as characterized by education, experience and learning ability, then rapid growth, in

[1] At Laitner and Stolyarov (2005) point out, this is a similar concept to the Solow (1960) idea of production capital retaining the technology level of the time it was produced in.

influencing these elements, should reduce the earning ability of the old compared to the young. Below is a (very simple) algebraic representation of this hypothesis:

$$\frac{w_y^H}{w_o^H} > \frac{w_y^L}{w_o^L} \tag{3.1}$$

where H denotes a high growth environment, L denotes low growth, y stands for 'young' and o for 'old'.

The effect of growth on the earning ability of the older cohorts can manifest through lower comparative wages and / or a higher rate of unemployment (increased difficulty in finding and retaining a job), e.g. an increase in job market ageism.

3.3 The Data: China Urban Household Survey, Years 1993-1997

My Data, acquired from the Universities Service Centre for China Studies in the Chinese University of Hong-Kong, contains cross-sectional data on roughly 31,000 Chinese households (approximately 100,000 individuals), over years 1993 to 1997 (about 6,200 households per year), spread over 57 cities, in ten of China's provinces. The different cities provide significant heterogeneity in average income, both cross-sectionally and over time.

Using this income heterogeneity, I am able to examine income and employment outcomes in both lower and higher growth environment. Since the data includes information on the income and employment of all household members, I can look both at comparative earnings and differences in employment uncertainty for younger and older cohorts.

3.4 A Look at the Data

First it is necessary to identify low and high growth environments. The time frame of my data, years 1993 to 1997, is insufficient to properly capture the long-term growth environment. I therefore proxy for growth using average 1997 income for each city. Since Chinese pre-reform society has been very egalitarian, current average income differences between locations should be highly correlated with the growth experienced in these locations. [2] In order to avoid issues pertaining to the

[2] Since both 1993 to 1997 growth and average 1997 income may be valid proxies, it is of note that they are 0.33 correlated, and that an OLS regression of 93-97 growth on average 97 income finds it statistically significant at the 5% level.

endogeneity of the family structure [3] I use only the income of the household head and spouse. I could alternatively measure average individual income in each city, but it the difference is negligible.

I break the Chinese cities in the data down to three groups: the general sample and high and low growth subsamples. The high and low growth subsamples are the top and bottom 25% of the sample in terms of growth, as proxied by the 1997 income level. Differences in average income between the groups are quite substantial: average household (head and spouse) income in the general sample is 10,010.82 Yuan, in the high sample it is 15,415.75 and in the low sample it is 6,546.15, so that the high sample is about 50% above the general sample and the low sample is about 50% below. Considering that the economic reform process that has created this inequality has started in the early 80's, these differences represent a substantial heterogeneity in growth rates.

Figure 3.1 below shows initial results of average household (head and spouse) income by age in low growth (bottom 25%) and high growth (top 25%) cities. Clearly, the income-age profiles are quite different, with average income in high growth cities reaching a plateau and starting to decrease much sooner in the lifecycle. This income also includes non-labor income such as pensions.

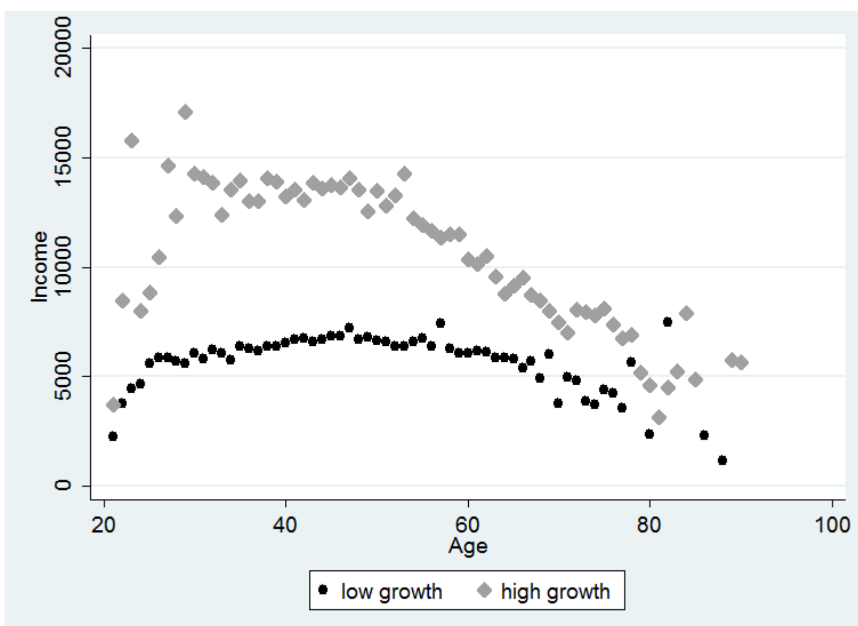

Figure 3.1: Household Level Income-Age Cross-section Under Low and High Growth

[3] The choice of having children and the choice of adult children to stay or leave the parent household may be influenced by growth and income levels.

Breaking away from the family structure and focusing on individual income, [4] we get similar results, as Figure 3.2 shows.

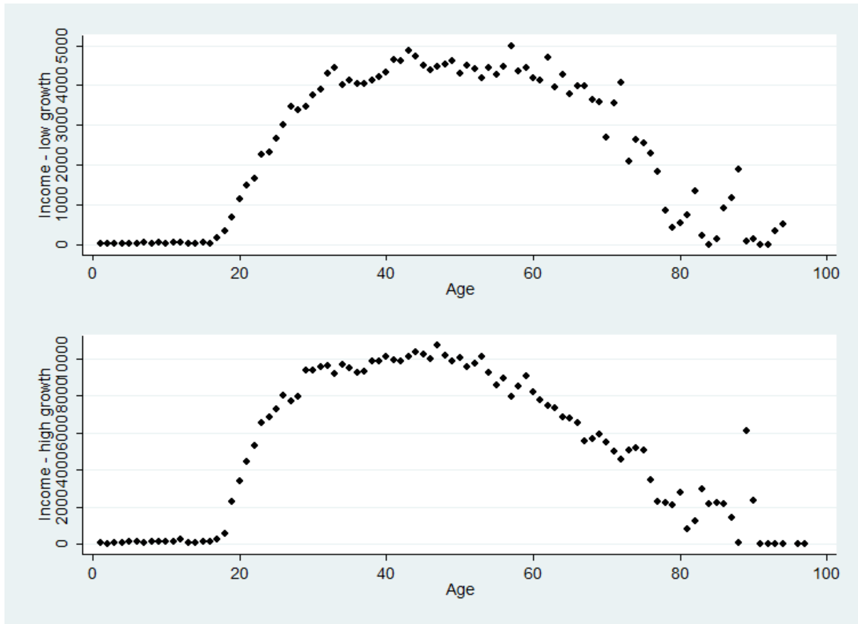

Figure 3.2: Individual Income-Age Cross-section Under Low and High Growth

Figures 3.1 and 3.2 deal with general income, including capital income, pensions and transfer payments. Figure 3.3 demonstrates the 'future shock' effect strictly for the labor income of the employed. Under high growth, wages of the young are relatively high at 20, increase faster up to around 30, and then reach a slower increase 'plateau'. The peak comes in the late 40's to early 50's, and wages seem to decline from there. Conversely, under low growth, the wages of the young are relatively low and seem to increase steadily over age until peaking in the late 50's.

Looking at employment, there does not seem to be a great deal of difference between the growth environments. Figure 3.4 shows the employment age profiles to be largely similar. If anything, the drop in employment around the age of 60 is sharper for the low growth regime.

[4] There may be endogeneity in being a household head, as elderly people without good income may choose to move in with their children.

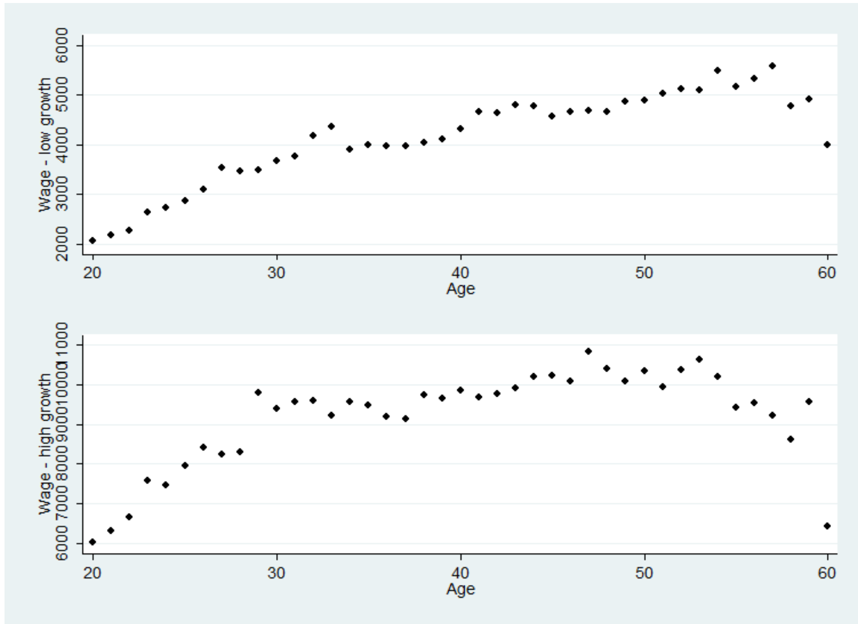

Figure 3.3: Individual Wage-Age Cross-section Under Low and High Growth

3.5 The Division Between State Owned and Private Employment

Many of the workers in the data (approximately 75%) are employed in China's state owned enterprises (SOE's). These firms are usually large scale and, as Song et al. (2011) note, with good access to credit. As such they may have less ability and motivation both to innovate and to track the productivity of individual workers, thus buffering their workers from the 'future shock' effect. The income age profiles in figures 3.5 and 3.6 seem to support this notion, with non-SOE workers showing a much larger effect. This result puts a different perpective on the claim, made in Song et al. (2011), linking the declining share of SOE's to the increase in household saving rates in China over the 1990's and 2000's.

Figure 3.7 shows the different wage paths of SOE and non-SOE workers. Clearly SOE employment is quite insulated from the 'future shock' effect. One can deduce that the reason for the flattening of the average income-age profiles over the last two decades in China, as witnessed in Chamon and Prasad (2010) and Song and Yang (2010), is the shift in employment shares between the state and private sectors, shown in Song et al. (2011).

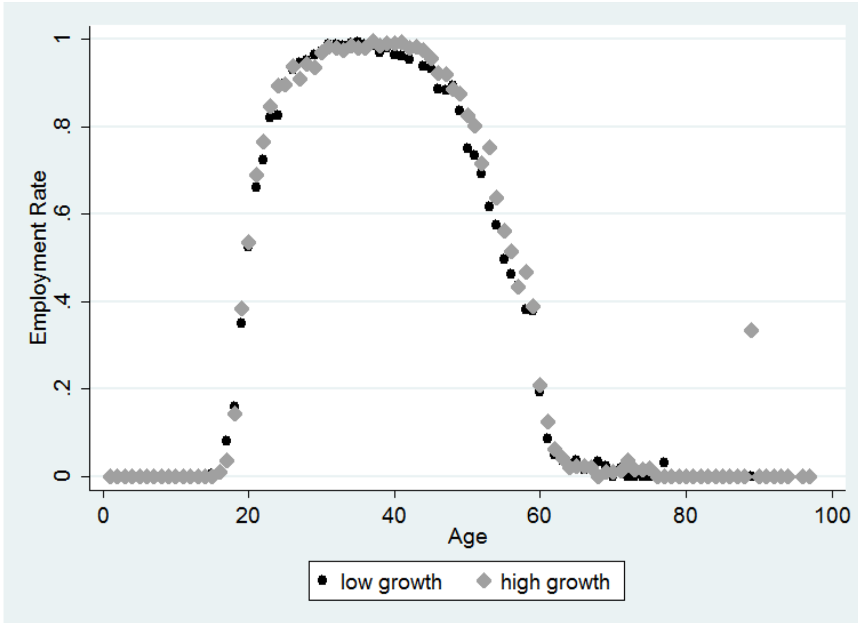

Figure 3.4: Employment-Age Cross-section Under Low and High Growth

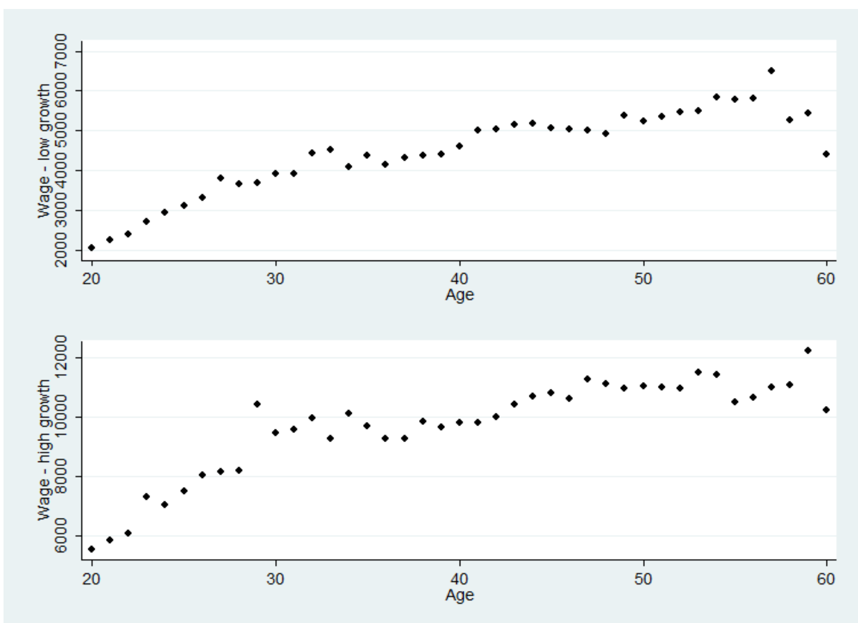

Figure 3.5: SOE Wage-Age Cross-section

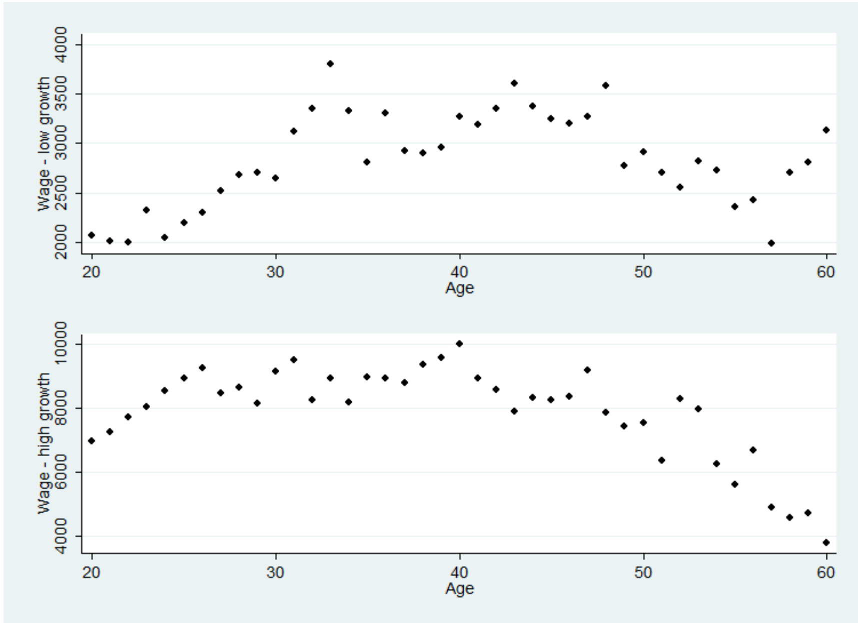

Figure 3.6: Non-SOE Wage-Age Cross-section

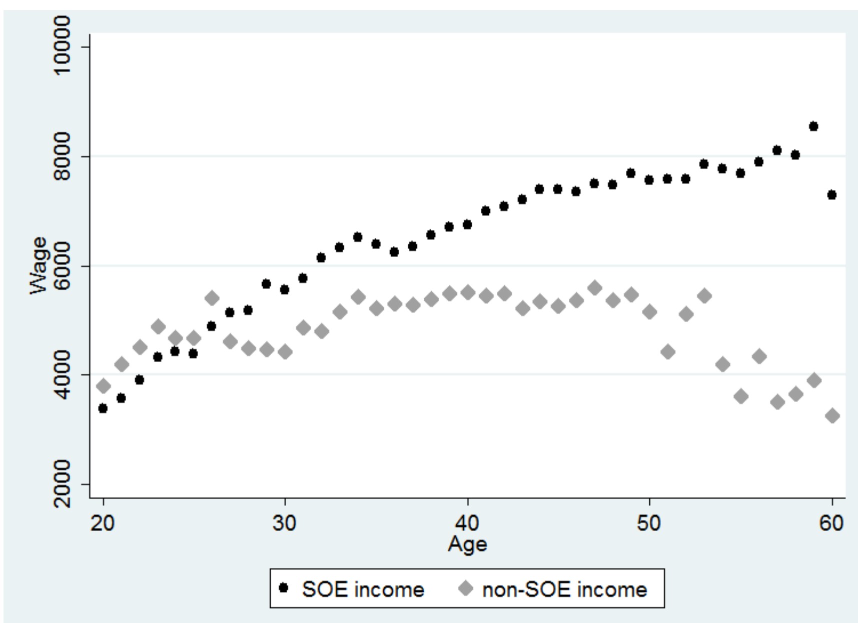

Figure 3.7: SOE vs. Non-SOE Wage-Age Cross-section

3.6 Estimating Future Shock

Table 3.1 presents the results of a Mincerian regression estimation using least squares with error terms clustered by city. On the left hand side is log of labor income normalized by average city income, and on the right hand side are age, experience, education and various dummy and interaction variables. Regressions are run over employed individuals, ages 20 to 60, with regression 1 including all of the sample, regression 2 including only SOE employees and regression 3 only non-SOE employees. Results show a strong 'future shock' effect. As expected, the effect is stronger in the non-SOE sectors on the economy.

log of normalized income	General	SOE	non-SOE
	(1)	(2)	(3)
high growth	.240	.164	.298
	(.070)***	(.068)**	(.096)***
low growth	-.053	-.049	-.079
	(.072)	(.050)	(.126)
age	.048	.050	.059
	(.007)***	(.009)***	(.013)***
age^2	-.002	-.003	-.002
	(.0004)***	(.0004)***	(.0008)***
age_H	-.026	-.012	-.040
	(.011)**	(.013)	(.014)***
age_L	.031	.022	.035
	(.010)***	(.012)*	(.018)*
age_H^2	.001	.0009	.001
	(.0006)*	(.0007)	(.0009)
age_L^2	-.001	-.0008	-.002
	(.0005)**	(.0006)	(.001)
xp	.021	.032	-.006
	(.004)***	(.005)***	(.005)
xp^2	.0001	-.0003	.0006
	(.00009)	(.00009)***	(.0001)***
xp_H	-.001	-.012	.008
	(.004)	(.008)	(.007)
xp_L	-.018	-.015	-.011
	(.007)**	(.006)**	(.010)
xp_H^2	.00002	.0002	-.0001
	(.0001)	(.0002)	(.0002)
xp_L^2	.0004	.0005	.0001
	(.0002)***	(.0001)***	(.0003)
high education	.367	.325	.322

log of normalized income	General	SOE	non-SOE
	(1)	(2)	(3)
	(.042)***	(.029)***	(.108)***
high school	.171	.152	.070
	(.032)***	(.023)***	(.058)
$high\,education_H$	-.153	-.110	.034
	(.057)***	(.055)**	(.129)
$high\,education_L$	.029	.025	.139
	(.059)	(.047)	(.183)
$high\,school_H$	-.089	-.035	.007
	(.047)*	(.046)	(.074)
$high\,school_L$	-.013	-.006	-.0005
	(.046)	(.041)	(.101)
age^3	.00003	.00004	1.00e-05
	(6.59e-06)***	(6.95e-06)***	(1.00e-05)
age^3_H	-.00002	-.00002	-1.00e-05
	(9.25e-06)*	(1.00e-05)	(1.00e-05)
age^3_L	1.00e-05	4.00e-06	.00003
	(8.99e-06)	(9.35e-06)	(.00002)
$high\,education * age$	-.0009	-.003	.003
	(.002)	(.002)**	(.005)
$high\,school * age$	.0006	-.001	.005
	(.001)	(.001)	(.003)*
$high\,education_H * age$	.007	.005	.002
	(.002)***	(.002)**	(.006)
$high\,education_L * age$	.001	.003	-.012
	(.003)	(.002)	(.009)
$high\,school_H * age$	.002	.0002	.0009
	(.002)	(.002)	(.003)
$high\,school_L * age$	.002	.002	-.006
	(.002)	(.002)	(.006)
1994	-.035	-.019	-.076
	(.006)***	(.007)***	(.017)***
1995	-.043	-.028	-.095
	(.006)***	(.008)***	(.019)***
1996	-.048	-.045	-.059
	(.007)***	(.008)***	(.021)***
1997	-.062	-.058	-.082
	(.008)***	(.009)***	(.022)***
Const.	-.876	-.868	-.870
	(.050)***	(.037)***	(.084)***
Obs.	55475	42906	12569

Table 3.1: Mincerian regressions over log of normalized income (clustered over 57 cities)

The 'future shock' effect is apparent in the higher initial income (the positive coefficient on the *high growth* parameter) and the flatter slope over age (the negative coefficient on age_H) compared to the general sample. These effect are noticeably stronger for the non-SOE sector. While the *low growth* coefficients are not statistically significant, they are all of the right sign. The slope of income over age appears to be steeper for the low growth sample (age_L is positive and significant). Here as well the effect is stronger for non-SOE sector.

The effect of growth on the return to experience are rather inconclusive, with xp_L and xp_H not systematically different from each other. It should be noted, however, that age and experience are highly (0.78) correlated, so that their effects may be intermingled. High education appears to have a lower return under high growth, with $high\,education_H$ negative in the general sample at a 1% significance level. Notice though that this effect is not significant for the non-SOE subsample.

3.7 Conclusion

Using growth heterogeneity between Chinese cities, I demonstrate a strong 'future shock' effect, shifting earning ability from older to younger cohorts. Breaking the effect down to education, experience and age elements, I show age to be the most important factor, with the caveat that high correlation between age and experience may cause the latter to appear less significant. Employment in itself does not appear to be affected, with people of the same age having a comparable employment ratio in high growth and low growth environments.

Dividing the sample to the SOE and non-SOE sectors, I show that the SOE sector is relatively shielded from the 'future shock' effect. Since, as Song et al. (2011) show, the share of the SOE sector has been falling quite rapidly in China over the past two decades, 'future shock' has had an increasing impact on society. This also explains the flattening out of the age-income profile in China over that same period. These results shed a different light on the welfare and distributional effects of growth, as well as on the link between growth and saving rates in a lifecycle context.

Bibliography

Aghion, Comin, and Howitt, "When Does Domestic Saving Matter for Economic Growth?," *NBER Working Paper*, 2006, (12275).

Andersson, Björn, "On the Causality Between Saving and Growth: Long- and Short-Run Dynamics and Country Heterogeneity," *Uppsala University Department of Economics Working Paper Series*, 1999, (18).

Ando, Albert and Arthur Kennickell, "How Much (or Little) Life Cycle is there in Micro Data? Cases of U.S. and Japan," in "Macroeconomics and Finance: Essays in Honor of Franco Modigliani," MIT Press, 1986.

Bakshi, Gurdip S. and Zhiwu Chen, "The Spirit of Capitalism and Stock-Market Prices," *American Economic Review*, 1996, *86* (1), 133–157.

Balassa, Bela, "The Purchasing Power Parity Doctrine: A Reappraisal," *Journal of Political Economy*, 1964, *72* (6), 584–596.

Baumol, William J., "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis," *American Economic Review*, 1967, *57* (3), 415–426.

Becker, Gary and H. Gregg Lewis, "Interaction between Quantity and Quality of Children," in T. W. Schultz, ed., *Economics of the Family*, 1973, pp. 81–90.

Carroll, Christopher D., "Why Do the Rich Save So Much?," in Joel B. Slemrod, ed., *Does Atlas shrug?: the economic consequences of taxing the rich*, Cambridge: Harvard University Press, 2000.

_ **and David N. Weil**, "Saving and Growth: A Reinterpretation," *Camegie-Rochester Conference Series on Public Policy*, 1994, *40*, 133–92.

_ **and Lawrence M. Summers**, "Consumption Growth Parallels Income Growth: Some New Evidence," in B. Douglas Bernheim and John B. Shoven, eds., *National saving and economic performance*, Chicago: Chicago University Press, 1991, pp. 305–43.

_ **, Jody Overland, and David N. Weil**, "Comparison Utility in a Growth Model," *Journal of Economic Growth*, 1997, *2* (4), 339–367.

_ **,** _ **, and** _ , "Saving and Growth with Habit Formation," *American Economic Review*, 2000, *90* (3), 341–355.

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

_ **, Kai Liu, and Eswar Prasad**, "Income Uncertainty and Household Savings in China," *IMF Working Paper*, 2010, (289).

Charles, Kerwin Kofi, Eric Hurst, and Nikolai Roussanov, "Conspicuous Consumption and Race," *The Quarterly Journal of Economics*, 2009, *124* (2), 425–467.

Christakis, NA and JH Fowler, "The collective dynamics of smoking in a large social network," *New England Journal of Medicine*, 2008, *358*, 2249–2258.

Christakis, Nicholas A. and James H. Fowler, "The spread of obesity in a large social network over 32 years," *New England Journal of Medicine*, 2007, *357* (4), 370–379.

Clark, Andrew E., Paul Frijters, and Michael A. Shields, "Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles," *Journal of Economic Literature*, 2008, *46* (1), 95–144.

Cole, Harold L., George J. Mailath, and Andrew Postlewaite, "Social Norms, Savings Behavior, and Growth," *The Journal of Political Economy*, 1992, *100* (6), 1092–1125.

Deaton, Angus, "Measuring Poverty in a Growth World (or Measuring Growth in a Poor World)," *The Review of Economics and Statistics*, 2005, *87* (1), 1–19.

_ **and Christina Paxson**, "Growth and Saving among Individuals and Households," *The Review of Economics and Statistics*, 2000, *82* (2), 212–225.

Duesenberry, James S., *Income, Saving and the Theory of Consumer Behavior*, Harvard University Press, Cambridge, MA, 1949.

Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes, "Do the Rich Save More?," *The Journal of Political Economy*, 2004, *112* (2), 397–444.

Easterlin, Richard A., "Does Economic Growth Improve the Human Lot? Some Empirical Evidence.," in R. David and M. Reder, eds., *Nations and Households in Economic Growth: Essays in Honor of Moses Abramovitz*, New York: Academic Press, 1974, pp. 89–125.

Francis, Johanna L., "Wealth and the capitalist spirit," *Journal of Macroeconomics*, 2009, *31* (3), 394–408.

Friedman, Milton, *A Theory of the Consumption Function*, Princeton, N.J.: Princeton University Press (for NBER), 1957.

Galí, Jordi, "Keeping up with the Joneses: Consumption Externalities, Portfolio Choice, and Asset Prices," *Journal of Money, Credit and Banking*, 1994, *26* (1), 1–8.

Goldberg, Jessica, "The Lesser of Two Evils: The Roles of Social Pressure and Impatience in Consumption Decisions," *Working Paper*, 2010.

Gourinchas, P.O. and J. Parker, "Consumption Over the Life Cycle," *Econometrica*, 2002, *70* (1), 47–89.

Harbaugh, Richmond, "Falling behind the Joneses: relative consumption and the growth-saving paradox," *Economic Letters*, 1996, *53*, 297–304.

I., A. Heston Kravis and R. Summers, "The Share of Services in Economic Growth," in Gerard Adams and Bert G. Hickman, eds., *Global Econometrics: Essays in Honor of Lawrence R. Klein*, MIT Press, 1983, pp. 188–218.

Jappelli, T. and M. Pagano, "Saving, Growth and Liquidity Constraints," *The Quarterly Journal of Economics*, 1994, *109* (1), 83–109.

King, Mervyn and L.D. Dicks-Mireaux, "Asset Holdings and the Life Cycle," *Economic Journal*, 1982, *92* (366), 247–67.

Kotlikoff, L.J. and L.H. Summers, "The Role of Intergenerational Transfers in Aggregate Capital Accumulation," *Journal of Political Economy*, 1981, *89* (41), 706–732.

Kuznets, Simon, *Shares of Upper Income Groups in Income and Savings*, New York: National Bureau of Economic Research, 1953.

Laitner, John and Dmitriy Stolyarov, "Technological Progress and Worker Productivity at Different Ages," *Michigan Retirement Research Center Working Paper*, 2005, (107).

Lee, Donghoon and Kenneth I. Wolpin, "Intersectoral Labor Mobility and the Growth of the Service Sector," *Econometrica*, 2006, *74* (1), 1–46.

Lee, R., S.H. Lee, and Andrew Mason, "Charting the economic life cycle," *NBER Working Paper*, 2006, (12379).

Luttmer, Erzo F.P., "Neighbors as Negatives: Relative Earnings and Well-Being," *The Quarterly Journal of Economics*, 2005, *120* (3), 963–1002.

Mariger, Randall P., "A Life-Cycle Consumption Model with Liquidity Constraints: Theory and Empirical Results," *Econometrica*, 1987, *55* (3), 533–557.

Modigliani, Franco, "The Life Cycle Hypothesis of Saving, The Demand for Wealth and the Supply of Capital," *Social Research*, 1966, *33* (2), 160.

_ , "The Life Cycle Hypothesis of Saving and Inter-Country Differences in the Saving Ratio," in W. A. Eltis et al., ed., *Induction, growth, and trade: Essays in honor of Sir Roy Harrod*, Oxford: Clarendon Press, 1970.

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

Ngai, Rachel and Christopher Pissarides, "Structural Change in a Multisector Model of Growth," *American Economic Review*, 2007, *97* (1), 429–443.

Obstfeld, M. and K. Rogoff, *The Foundations of International Macroeconomics* 1996.

Oshio, Takashi, Kayo Nozaki, and Miki Kobayashi, "Relative Income and Happiness in Asia: Evidence from Nationwide Surveys in China, Japan, and Korea," *Social Indicators Research*, 2010, pp. 1–17.

Parente, S.L. and E.C. Prescott, "Changes in the Wealth of Nations," *Federal Reserve Bank of Minneapolis Quarterly Review*, 1993.

Rosenquist, JN, J Murabito, JH Fowler, and NA Christakis, "The spread of alcohol consumption behavior in a large social network," *Annals of Internal Medicine*, 2010, *152* (7), 426–33.

Samuelson, P. A., "Theoretical notes on trade problems," *Review of Economics and Statistics*, 1964, *46*, 145–154.

Solow, Robert M., "Investment and Technological Progress," in Kenneth Arrow, Samuel Karlin, and Patrick Suppes, eds., *Mathematical Methods in the Social Sciences*, Stanford University Press, 1960.

Song, Zheng Michael and Dennis Tao Yang, "Life Cycle Earnings and the Household Saving Puzzle in a Fast-Growing Economy," *Working paper*, 2010.

_ , Kjetil Storesletten, and Fabrizio Zilibotti, "Growing Like China," *American Economic Review (forthcoming)*, 2011.

Tobin, James, "Life Cycle Saving and Balanced Growth," in "Ten Economic Studies in the Tradition of Irving Fisher," J. Wiley, 1967.

UN, "Living arrangements of older persons around the world," UN Department of Economic and Social Affairs Population Division 2005.

Veblen, Thorstein, *The Theory of the Leisure Class*, McMillan, 1899.

Weber, Max, *The Protestant Ethic and the Spirit of Capitalism*, 1958 translation, Charles Scribner and Sons, New York, 1905.

Wei, Shang-Jin and Xiaobo Zhang, "The Competitive Saving Motive: Evidence From Rising Sex Ratios and Saving Rates in China," *NBER Working Paper*, 2009, (15093).