A Relative Answer to the Growth-Saving Puzzle

Gruber, Noam

6 October 2020

A Relative Answer to the Growth–Saving Puzzle *

Noam Gruber †

October 6, 2020

Abstract

Prolonged rapid growth, i.e. a "catching-up" process, is known to be accompanied by high rates of household saving. This phenomenon is central in explaining the direction of international capital flows and trade imbalances in the past several decades, yet it is very much in contradiction to prevailing macroeconomic theory. This paper finds that a standard life-cycle model, even when integrated with uncertainty about future growth and credit constraints, is completely unable to replicate the relations between growth and saving as seen in the data. However, adding utility from relative consumption to the model allows for the full replication of these relations.

Keywords: Life Cycle, Saving Growth, Open Economy Growth, Household Saving, Life Cycle Models and Saving, Relative Income Hypothesis

JEL Classification Numbers: D91, E21, F43, O11.

*I am grateful to Alan Deardorf, John Laitner, Jagadeesh Sivadasan and Dmitriy Stolyarov for their helpful advice. I thank participants of the Midwest International Economic Development Conference at the University of Minnesota, the Midwest Macroeconomic Meetings at the University of Colorado Boulder and seminar participants at Dartmouth College and the University of Michigan for their comments. Opinions expressed here are solely my own and do not represent the views of the National Economic Council of Israel. All errors are my own.

†National Economic Council of Israel. Email: grubern@pmo.gov.il.

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.

However, 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 slow-growing rich countries at the expense of its future wealth. Reality is often the complete opposite.

It has been suggested that it is the higher propensity to save that is causing the higher growth, a-là a Solow style model. However, as Carroll and Weil (1994) and Andersson (1999) show, it is higher growth that precedes higher saving rates and not the other way around; hence it is more likely that causality runs from growth to saving. While a positive causality from saving to growth cannot be ruled out,[1] clearly the saving decision must be optimal given growth expectations. Since much of the greater saving-to-GDP ratio is attributed to higher household saving rates,[2], than the question of saving-rate optimality should be discussed from the household's perspective, rather than some central planner's.

Another suggested explanation is that both saving and growth are attributable to some cultural characteristics. For example, the rapid growth and high saving rates in East Asia can be the result of Confucian values shared by many countries in that region. This line of argument fails, however, to account for several facts. First, high saving rates have historically accompanied prolonged rapid growth (a "catching-up" process) in many countries across the

[1] For example, Aghion et al. (2016) suggest a model in which domestic savings allow domestic banks to co-finance foreign technologically advanced investment and thus foster growth.

[2] For example, Ando and Kennickell (1986) find that during the high-growth period in Japan, households in all age groups there were saving much more than their US counterparts. Chamon and Prasad (2010) show a rapid increase in the Chinese household saving rate over the last several decades.

world. While much of the current discussion is about China, this phenomenon has been seen in post world-war II Europe, in Japan, the four tigers (Korea, Taiwan, Hong-Kong and Singapore), in Ireland, Chile, and more recently in Vietnam and India. Second, as discussed above, saving increases following a rise in growth. Why then have Confucian values not made households save in high rates all along? Finally, Carroll et al. (1994) and Carroll et al. (1999) look at the effect of cultural background on saving behavior among immigrants, and do not find a correlation with the saving rates at the countries of origin. Therefore, it seems appropriate to investigate the growth-saving correlation in generality, and not to relegate it to cultural attributes.

Modigliani (1966) attempts to explain the causality from growth to saving using a life-cycle model. The intuition behind this explanation is simple: 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.[3] 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"), examining Chinese data, Chamon and Prasad (2010) find no evidence of it. Gruber (2018) demonstrates empirically that household saving behavior is determined by relative, rather than absolute, income level, and argues that utility from

[3]The latter paper also contains an illuminating literature review of the topic.

relative income and wealth can explain both the higher saving rates by the relatively rich and the positive relation between growth and saving.

In addition to its effect on the saving rate, rapid growth seems to have two more effects, closely related to each other: 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. (2014), which Carroll et al. (2000) seek to explain with habit formation, is also inconsistent with the standard Permanent Income Hypothesis (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 quantifies and measures the ability of a life-cycle model to generate the following three stylized facts, observed in the literature: one, a positive relation from aggregate growth to the aggregate household saving ratio (with growth fully affecting the income path of individual cohorts); two, a correlation between the aggregate growth rate and the household expenditure growth rate over the life cycle; and three, a consumption-age cross-sectional distribution that is largely invariant to aggregate growth.

Furthermore, this paper looks at the effect of augmenting the life-cycle model with three elements - uncertainty about future growth, credit constraints and utility from relative consumption (also known as “keeping up with the Jonses”) - on its ability to generate the stylized facts. While the standard elements – growth uncertainty and credit constraints – are insufficient in accounting for the growth-saving relation, the element of utility from relative consumption allows the model to replicate all three stylized facts: the correlation between growth and saving, the rapid increase of consumption expenditure over the household life-cycle and the stability of the consumption-age cross section, as observed in growing economies.

Section 2 presents the life-cycle model used, and the addition of growth uncertainty and credit constraints to the model. Section 3 describes the modeling of utility from relative consumption. The respective foci of sections 4 and 5 are the calibration and simulation of the model. Section 6 concludes.

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. 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:

$$U = \sum_{j=1}^{J} \beta^{j-1} u\left(c_j, j\right), \tag{1}$$

where c_j, abbreviated from $c(j, t, g, r)$ for convenience, is the consumption of a household at age j at time t, under growth rate g and interest rate r. The household's intra-period utility varies with age, and is of the following form:

$$u\left(j, c_j\right) = q\left(j\right) \frac{c_j^{1-\theta} - 1}{1 - \theta}, \tag{2}$$

where $q(j)$ is what shall be referred to as the "life-cycle multiplier", putting weights 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] Here the path of the life-cycle multiplier over the life of the household will be referred to as the "household regime" (i.e. $\{q(j)\}_{j=1}^{J}$).

The demographic transition that typically takes place over the development process may also affect the households' saving decisions. One could argue, for example, that since we observe the average number of children per household diminishes with economic moderniza-

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

tion, the household regime should vary with the level of development. In addition, changes in birth rates would also lead to changes in the elderly support ratio. Another example 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 that prolonged development, such as discussed in this paper, lowers the rate in which this phenomenon takes place.[5] Thus a young person expecting long term growth should not overly rely on the support of his or her children after retirement. Curtis et al. (2015) argue that these demographics changes play a large role in explaining the increase in the Chinese household saving rate. While that may be the case in China, the phenomenon of high growth and high saving rates has occurred in different countries with different demographic compositions. Therefore, this paper abstracts away from this element. In light of the above, and following evidence (discussed in Section 4) that the consumption-age profile is very stable, the household regime is assumed 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. The total number of effective units of labor at the economy is assumed to be constant at unity:

$$\sum_{j=1}^{J} l(j) = 1. \tag{3}$$

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, $W(t,g)$:

$$w(j,t,g) = l(j)W(t,g), \tag{4}$$

where $w(j,t,g)$ is the labor income of a household of age j at time t under a prevailing constant growth rate of g. $W(t,g)$ is growing at rate g, hence growth fully affects the

[5] See UN (2005).

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,[6] so that the interest rate, r, is exogenous, discounted lifetime income is:

$$\Omega(t,g,r) = \sum_{j=1}^{J} (1+r)^{-j+1} w(j,t+j-1,g) = \sum_{j=1}^{J} (1+r)^{-j+1} l(j)W(t+j-1,g). \quad (5)$$

The discounted lifetime consumption expenditure of a household created at time t in an economy with a constant growth rate of g and with an international interest rate r is:

$$C(t,g,r) = \sum_{j=1}^{J} (1+r)^{-j+1} c_j. \quad (6)$$

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

$$\begin{aligned} &\max \sum_{j=1}^{J} \beta^{j-1} q(j) \frac{c_j^{1-\theta}-1}{1-\theta} \\ &s.t. \\ &C(t,g,r) = \Omega(t,g,r), \end{aligned} \quad (7)$$

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

2.1 Uncertainty of Future Growth

In the general life-cycle model above the households are assumed to 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

[6] This 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.

world can expect are relatively mild, that may not be the case for a household in a rapidly developing country.[7] Such a country is in a process of catching up to 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 quite difficult to establish the point in which a country has reached its development frontier.[8]

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

[7]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.

[8]For an in depth modeling and discussion of the catching-up process, see Acemoglu et al. (2006).

in the following fashion: Each period, all households perceive a risk of 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).[9] 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:

$$V^H(a_j, j) = \max_{c_j} \quad u(c_j, j) + \beta \left\{ \lambda V^L(a_{j+1}, j+1) + (1-\lambda) V^H(a_{j+1}, j+1) \right\} \tag{8}$$

$$s.t. \quad a_{j+1} = (a_j + w_j - c_j)(1+r), \quad a_{J+1} = 0, \tag{9}$$

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

$$V^L(a_j, j) = \max_{c_t} \quad \sum_{t=j}^{J} \beta^{t-j} u(c_t, t), \tag{10}$$

$$s.t. \quad \sum_{t=j}^{J} (1+r)^{j-t} c_t = a_j + \sum_{t=j}^{J} (1+r)^{j-t} l(t) W(j) (1+g_L)^{t-j}. \tag{11}$$

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 8 to 10 lead to the following Euler equation:

$$u_c(c_j^H, j) = (1+r)\beta \left\{ \frac{\lambda u_c(c_{j+1}^L, j+1)}{(1+\phi_L)^{1-\gamma}} + \frac{(1-\lambda) u_c(c_{j+1}^H, j+1)}{(1+\phi_H)^{1-\gamma}} \right\}, \tag{12}$$

[9] It is reasonable to believe that λ should actually become higher as the country approaches the development frontier. Such variations in λ are abstracted from for simplicity.

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 $j+1$.

2.2 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. henceforth referred 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.

This paper uses the following credit constraint formulation:

$$a_j \geq \kappa \Omega_j, \tag{13}$$

where Ω_j is expected discounted future labor income, so that:

$$\Omega_j = \sum_{t=j}^{J} (1+r)^{j-t} E_j \left[w_t \right], \tag{14}$$

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 as 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.

The algorithm detailed in Mariger (1987) is used to compute consumption paths given credit constraints.

3 Relative Consumption

The role of relative consumption has long been a topic of economic discussion, as evidenced by Veblen (1899) and Duesenberry (1949). In a finding known as the 'Easterlin paradox', Easterlin (1974) observes that while the rich are generally happier than the poor within a given country, average happiness in rich countries is not greater than in poor ones, and economic growth does not seem to increase average happiness over time. More recently, Luttmer (2005) finds that higher levels of consumption by neighbors lead to lower reported household happiness levels. Jin et al. (2011) find that local inequality increases household saving rates in China, and Gruber (2018) establishes relative income to be a central determinant of household saving behavior, also for China.

Consider the following function, based on the Galí (1994) formulation, for the household's within-period utility:[10]

$$u(j, c_j, \bar{c}_{t,g}) = q\,(j)\,\frac{c_j^{1-\theta}\bar{c}_{t,g}^{\gamma\theta}}{1-\theta}. \tag{15}$$

c_j, as noted in Section 2, is the consumption of a household of age j at time t, given growth rate g and a constant international interest rate r. $\bar{c}_{t,g}$ is the average household consumption at time t given growth rate g, and γ is assumed to be between zero and one. $q(j)$ is the household regime, as detailed in Section 2.

The lifetime utility function is therefore:

$$U = \sum_{j=1}^{J} \beta^{j-1} u(j, c_j, \bar{c}_{t,g}), \tag{16}$$

The above, in conjunction with a standard budget constraint (see Equation 7), yields the following Euler equation:

$$c_{t+1} = \left[\frac{q(t+1)}{q(t)}(1+r)\beta\right]^{\frac{1}{\theta}} (1+g)^{\gamma} c_t \tag{17}$$

As Harbaugh (1996), Carroll et al. (1997) and Gruber (2018) discuss, a relative consumption element may generate a causal relation from economic growth to the saving rate. Equation 17 demonstrates that 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 on.

[10]This formulation is chosen for its elegance. However, other formulations could generate similar dynamics. See for example Brunnermeier and Nagel (2008).

4 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 (i.e. $q(j)$ and $l(j)$). For the former, values within the commonly accepted ranges are used, as follows:

Parameter	**Value**
T (length of household lifecycle)	55
θ (consumption utility risk averseness coefficient)	4
α (Cobb-Douglas capital intensity)	0.33
β (time preference)	0.95
δ (depreciation rate)	0.05
r (international real interest rate)	5%

Table 1: Calibration of Model Parameters

The life cycle parameters $q(j)$ and $l(j)$ are calibrated using US data. The reason for using US data, rather than data from a rapidly growing country, is to avoid the suspicion that there is something special about such a country's income and/or consumption age profile that would generate an saving behavior different than that of other countries. The 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).[11]

Using the fact that $c(j,t,g,r) = c(j,(t-j+1)+j-1,g,r)$, $q(j)$, the household's life-cycle multiplier, can be backed out of actual household consumption data, assuming that $1+r \approx \frac{1}{\beta}$. [12] First, the equation is rewritten in matrix form:

[11] This assumption may seem somewhat problematic, considering the changes in longevity and social structure that come hand in hand with economic development. However, it is still 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 much smaller.

[12] A small open economy assumption.

$$C_t = X \cdot q/(Y_t q), \tag{18}$$

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, $X(j) = (1+g)^{-j+1}$, and $Y_t(i,j) = (1+r)^{-j+1}$.

In other words, the model attributes consumption variation across cohorts to the different values of $q(j)$ as well as growth.[13] Equation 18 above can be rewritten in the following form:

$$C_t/X\iota' \cdot Y_t q = q, \tag{19}$$

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

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

US Consumer Expenditure (CEX) survey data on household consumption for years 1984-2005 is used.[14] As can be seen in figure 4, consumption patterns in the data remain very consistent over the years.[15] 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 is very similar to the findings of Gourinchas and Parker (2002) and Fernández-Villaverde and Krueger (2007).

Having experimented with a variety of household regimes, I find that, save for extreme and unreasonable cases, using different household regimes of the general hump shape does

[13]It is worth noting that other elements, (i.e. growth uncertainty, credit constraints and utility from relative consumption), 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 in Section 5, these elements should not have a substantial influence.

[14]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.

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

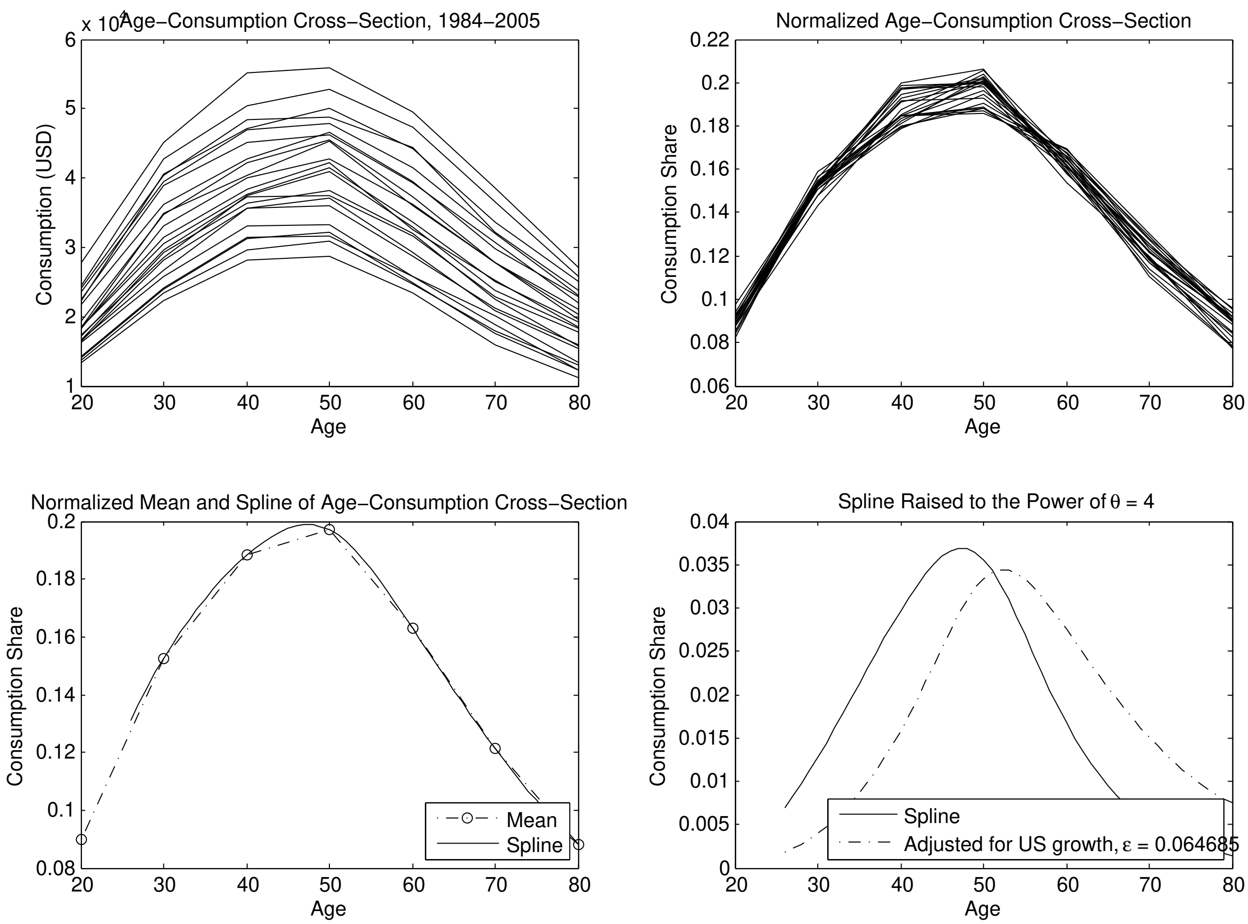

Figure 1: Calibrating the household regime, the weighting (q) on the household consumption utility by period

not affect results substantially.

Calibrating $l(j)$, the effective labor units of the household over its life-cycle, is even more straight forward: again use CEX data (years 1984-2005) on the wage income of different household cohorts is used 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 assumes that the distribution of effective 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 income for the last age bracket, since the number of observations for that category drops considerably, possibly causing a selection bias.

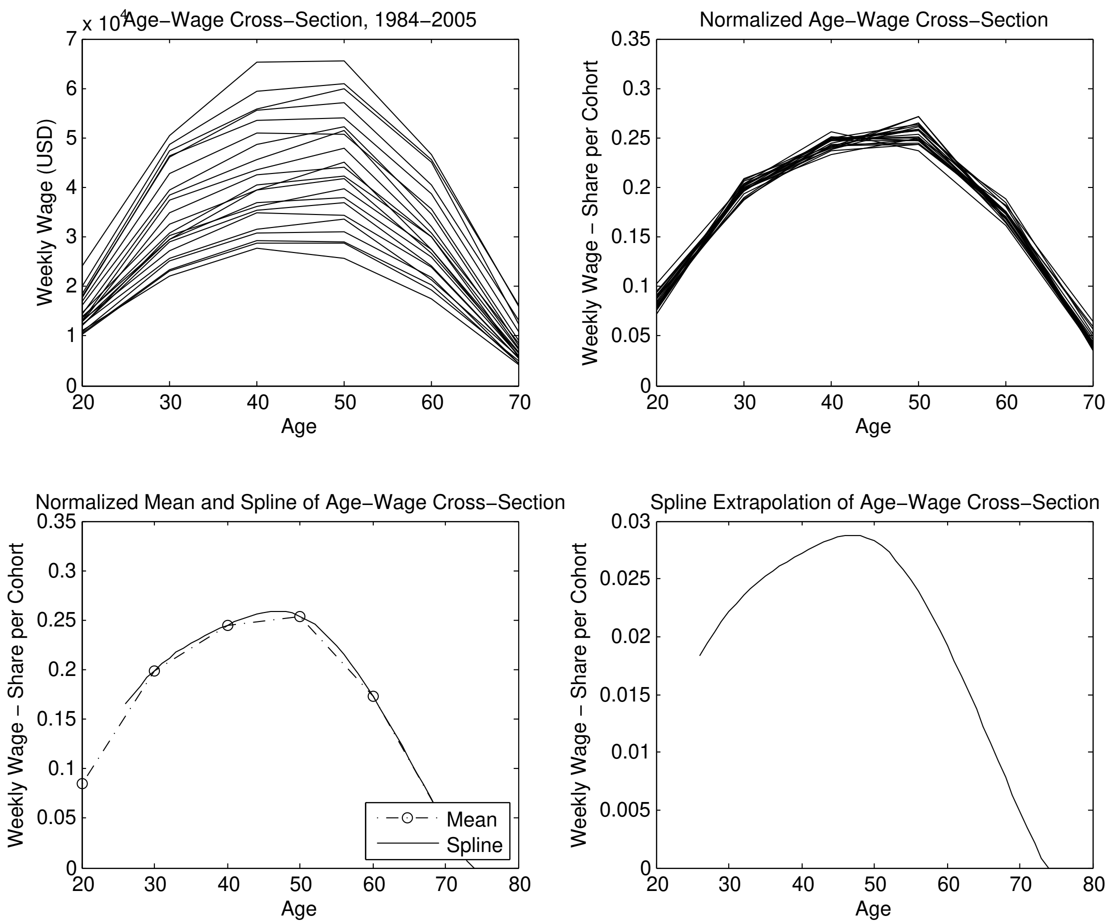

Figure 2: The household's effective labor units (l) over the life cycle

5 Simulation

5.1 The Life Cycle Baseline Model

The baseline model has constant growth with 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 monotonically positive correlation between growth and saving seen in the data.

In conjunction with the investment line, 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.

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

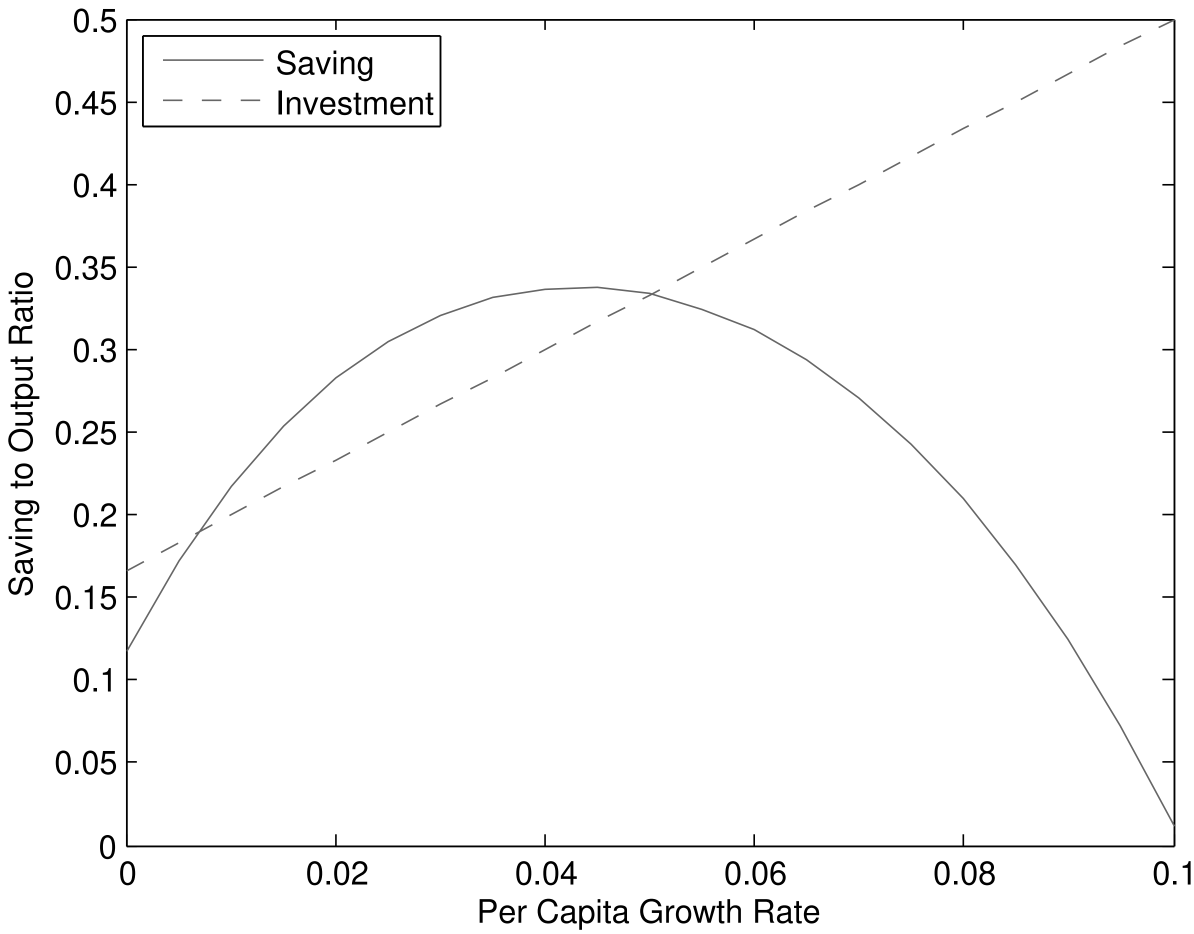

Figure 3: Saving and investment ratios as a function of the growth rate, baseline life cycle model

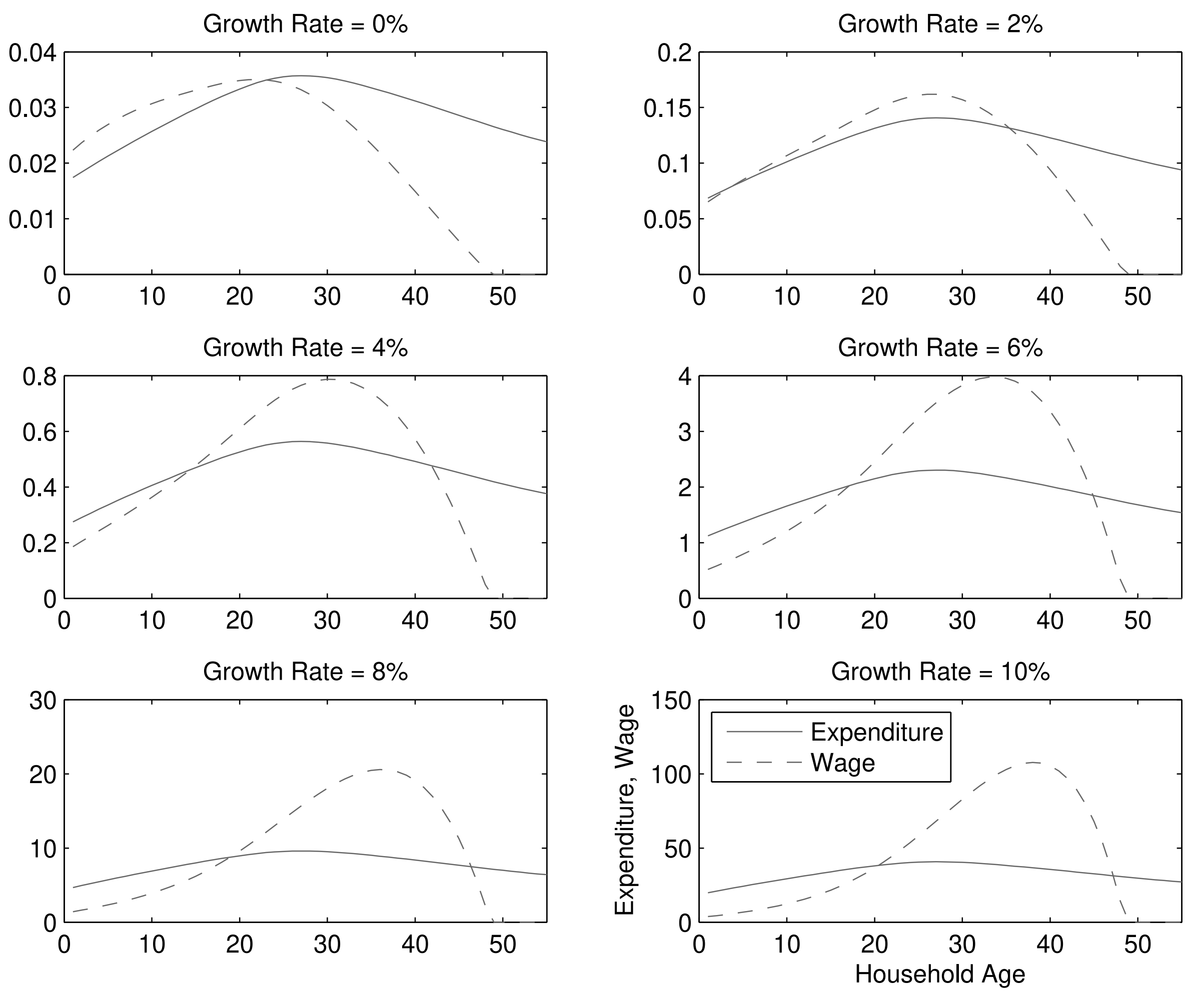

Figure 4: Household life cycle expenditure and wage, baseline model

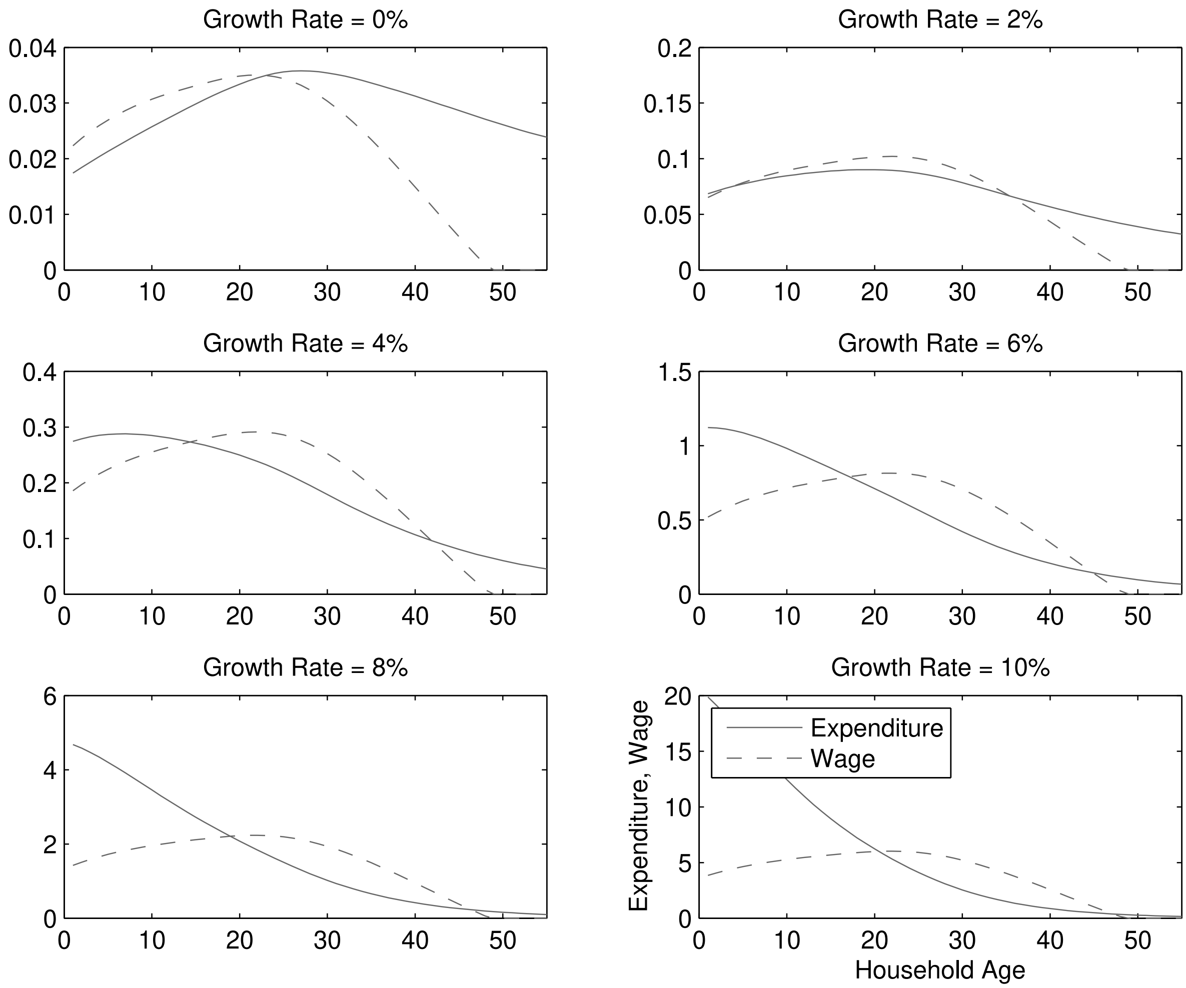

Figure 5: Cross section of household expenditure and wage, baseline model

correlated and higher aggregate growth is demonstrated in both Carroll and Summers (1991) and Lee et al. (2014).

Another 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. (2014) 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.

5.2 Growth Uncertainty

Growth uncertainty is modeled here as a constant probability λ that the economy will drop from its high growth state to an absorbing low growth state (the high growth state cannot be resumed). 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 6 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 7 shows, though at higher growth rates than the baseline case, as shown in Figure 8.

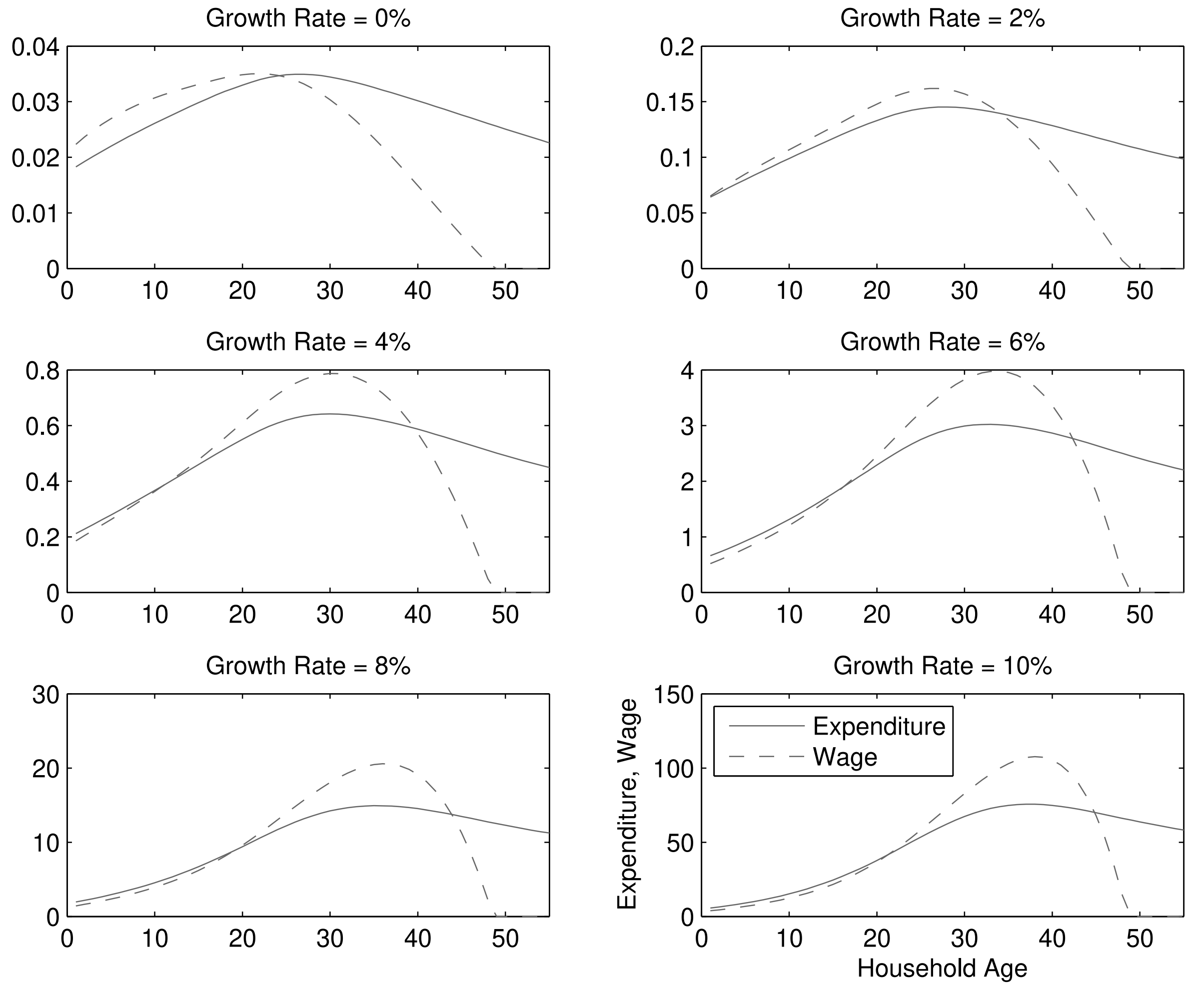

Figure 6: Household life cycle expenditure and wage, with growth uncertainty ($\lambda = 5\%$)

Notice also that comparing Figures 4 and 6, 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 major bias.

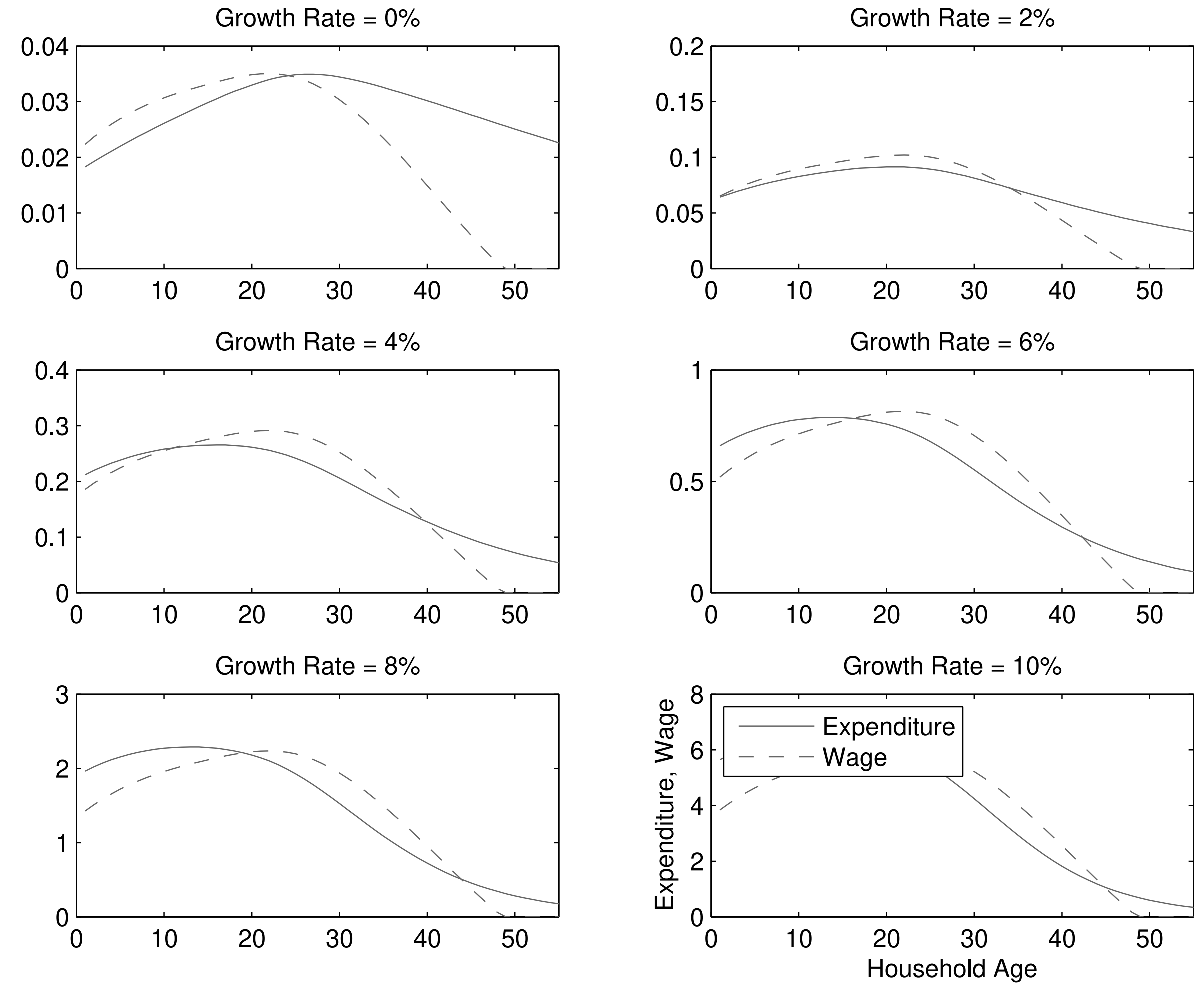

Figure 7: Cross section of household expenditure and wage, with growth uncertainty ($\lambda = 5\%$)

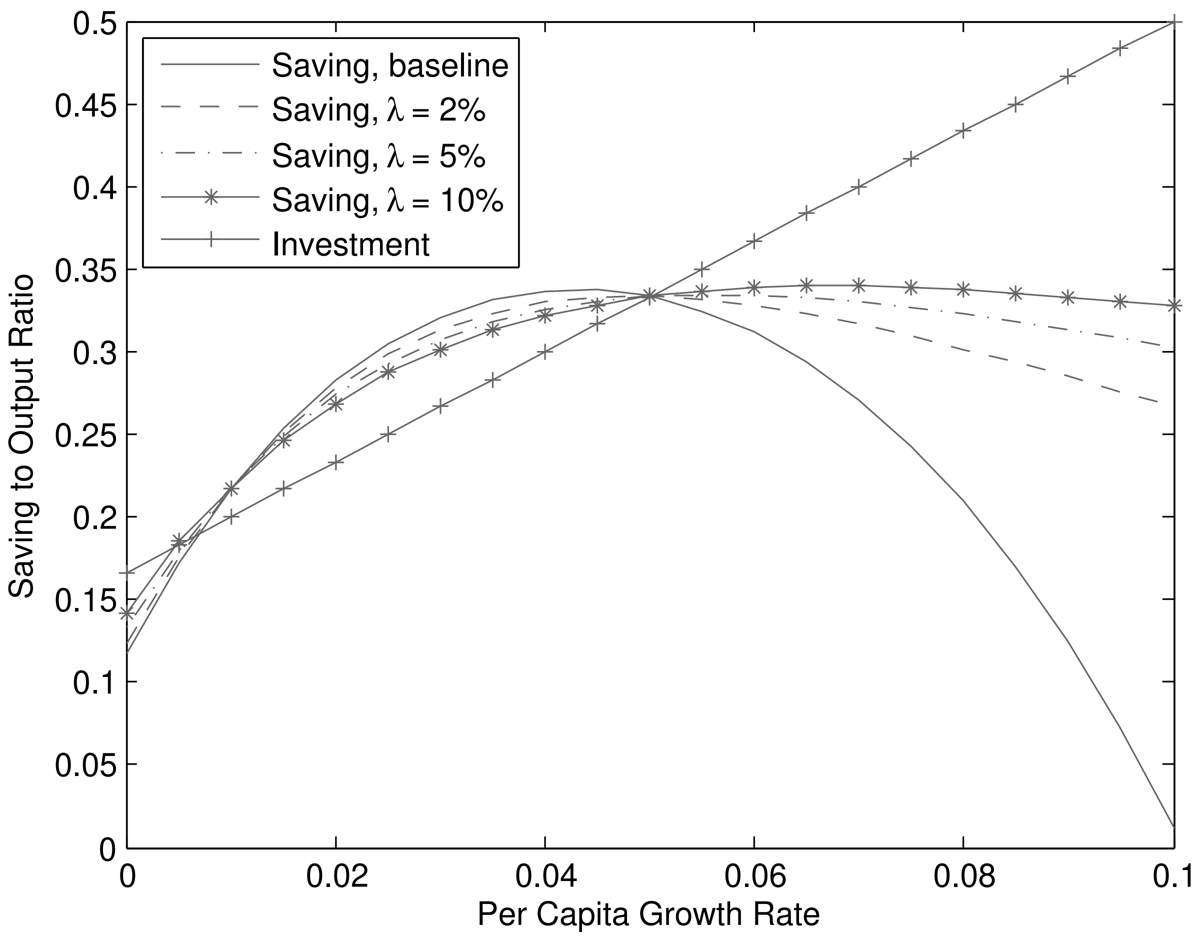

Figure 8: Saving and growth - a sensitivity analysis for different levels of growth uncertainty

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 8), 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-cycle expenditure profiles.

In summary, adding growth uncertainty does not allow the life-cycle model to replicate the three stylized facts of growth and saving.

5.3 Credit Constraints

Credit constraints, inhibiting the ability of younger generations to dissave, can have a powerful impact on aggregate saving. As Figure 9 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 9 to stop the negative effect of growth on the aggregate saving rate seen in the standard life-cycle model above 4% per capita growth, and creates a saving rate plateau all the way up to 10% per capita growth (the upper limit of the simulation).

Figure 10 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 the

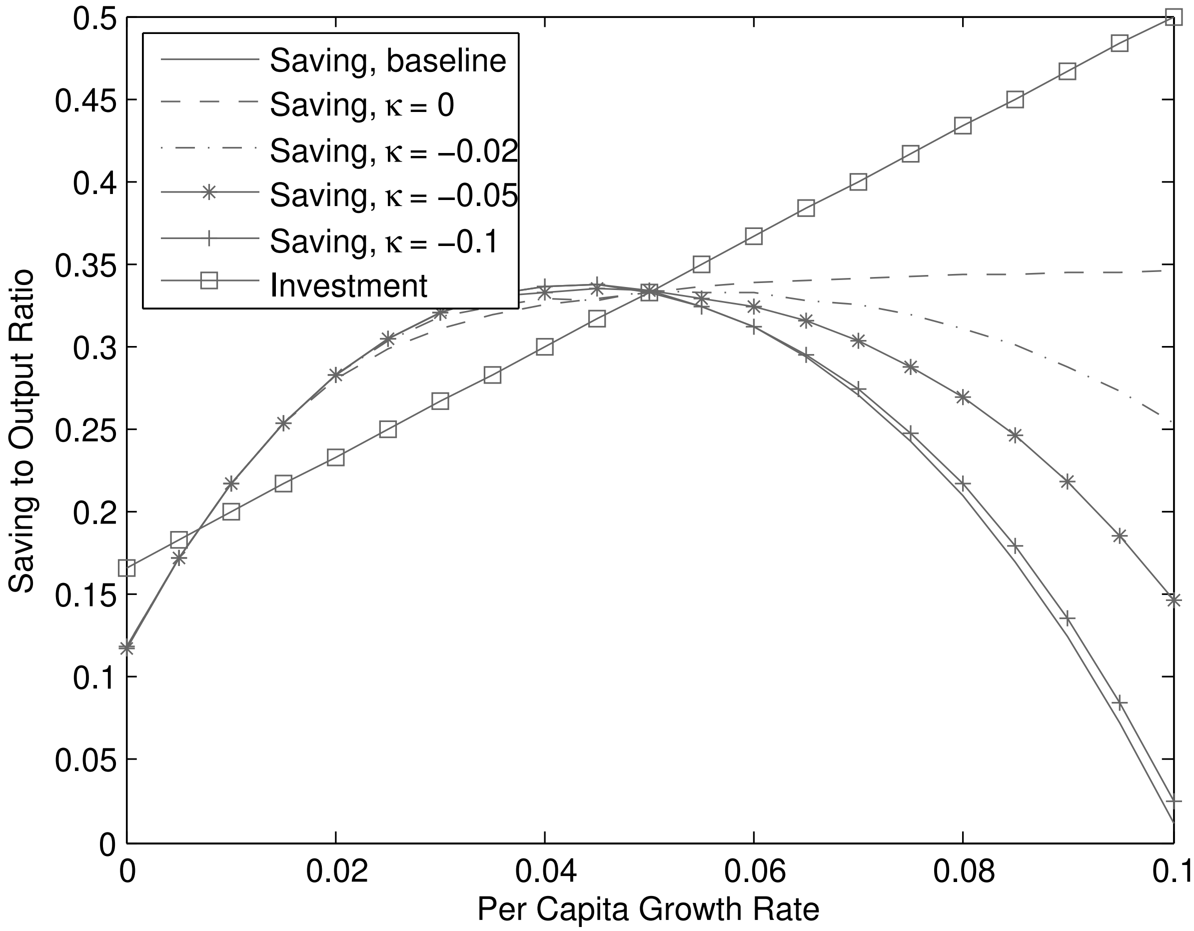

Figure 9: Saving and growth - a sensitivity analysis for different levels of credit constraints

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.

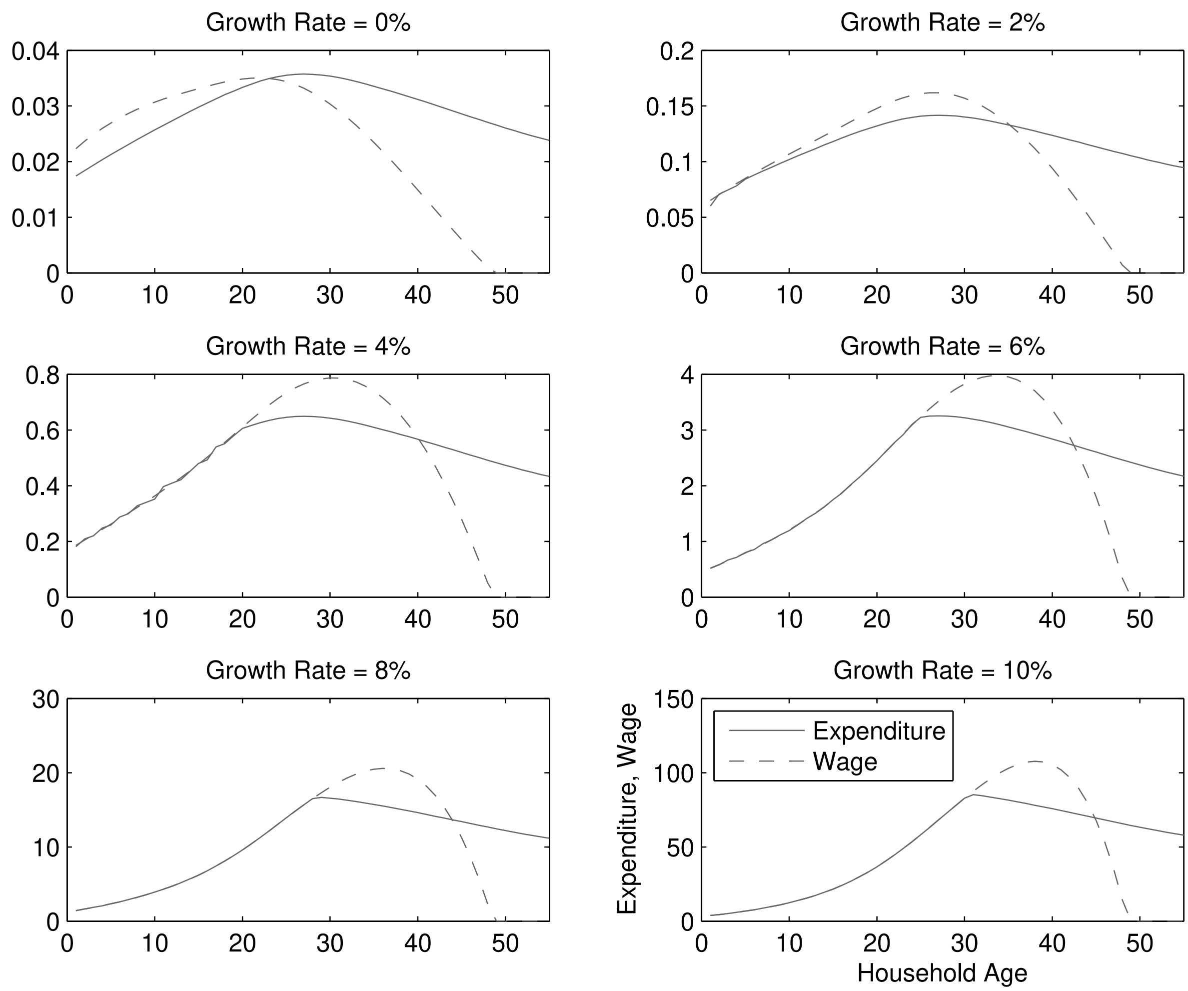

Figure 10: Household life cycle expenditure and wage, with credit constraints ($\kappa = 0\%$)

As Figure 11 shows, credit constraints also 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.

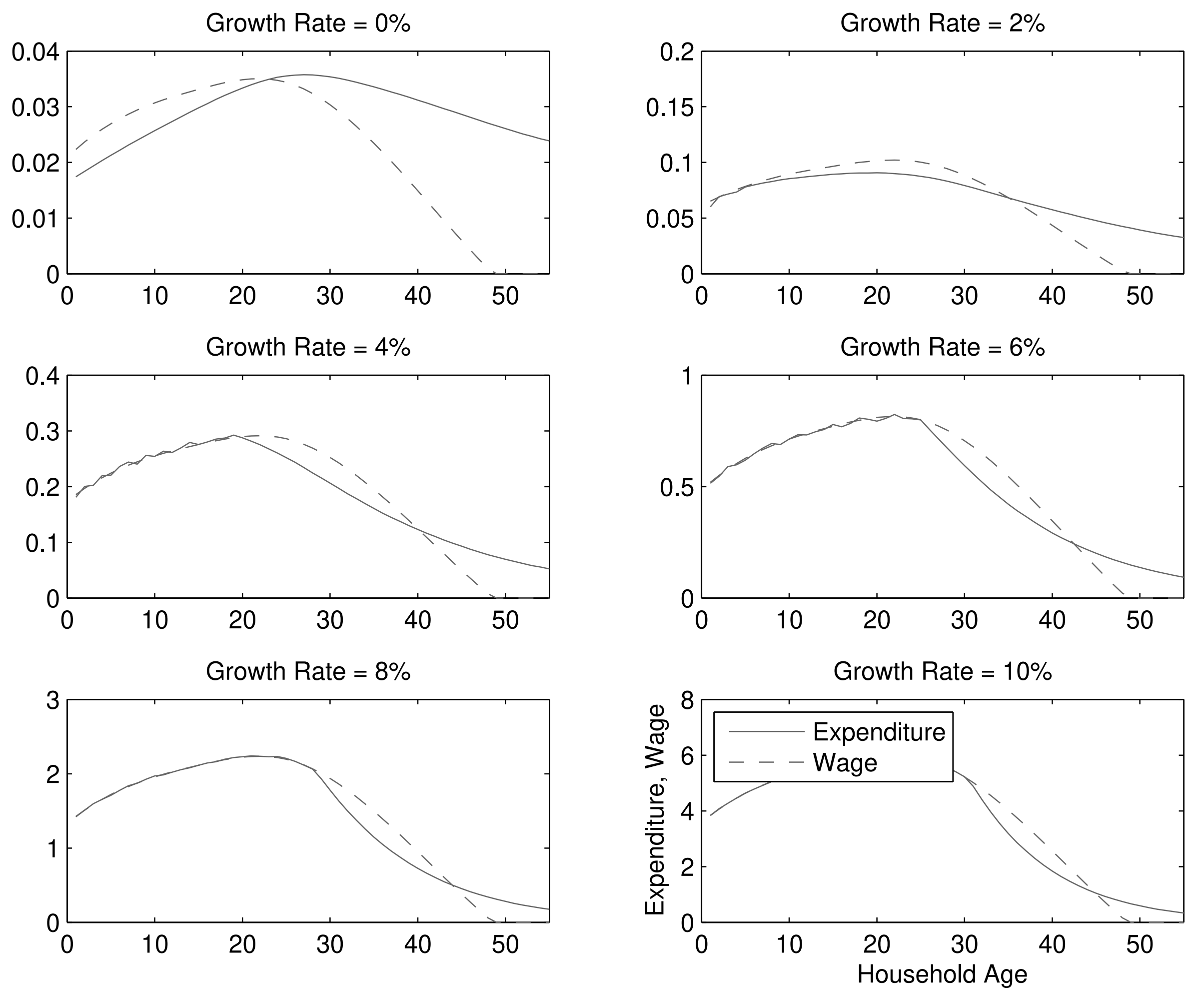

Figure 11: Cross section of household expenditure and wage, with credit constraints ($\kappa = 0\%$)

In summary, adding growth uncertainty does not allow the life-cycle model to replicate the three stylized facts of growth and saving.

5.4 Relative Consumption

With γ set to one (see Equation 15), the effects of utility from relative consumption are simulated. As can be seen in Figure 12, the increase in average consumption, caused by economic growth, does indeed lead the households to consume a larger proportion of their income during later years, and to continue to save when young despite their increasing income path. Young cohorts not only save instead of dissaving, but their period of saving is actually extended to later years in the lifecycle. Utility from relative consumption thus enables the lifecycle model to replicate the correlation between the aggregate growth rate and the growth rate of individual consumption over the lifecycle.

The effect is strong, and substantially changes the aggregate picture, as Figures 13 and 14 demonstrate. Since growth pushes young cohorts to save rather than dissave, than, just as Modigliani (1970) argues, the higher lifetime wealth of the young compared to the old increases the aggregate saving rate. In addition, as Figure 13 shows, growth does not skew the consumption-age cross section.

Figure 14 shows the effect of growth on the aggregate saving rate with utility from relative consumption ($\gamma = 1$). The effect is different than under the growth uncertainty and credit constraint elements, as the positive effect of growth on saving now extends to higher growth levels instead of tapering out. Notice also that the implication of the results presented in Figure 14 is that slow growth countries should have trade deficits, whereas high growth countries would have surpluses, with a cut-off growth rate of about 5% per capita. It is important to note in this context that this is a partial equilibrium analysis, holding the interest rate constant. Different interest rates would change both the investment and saving rates for the different growth levels, yielding different cut-offs.

These results confirm that utility from relative consumption allows the life-cycle model

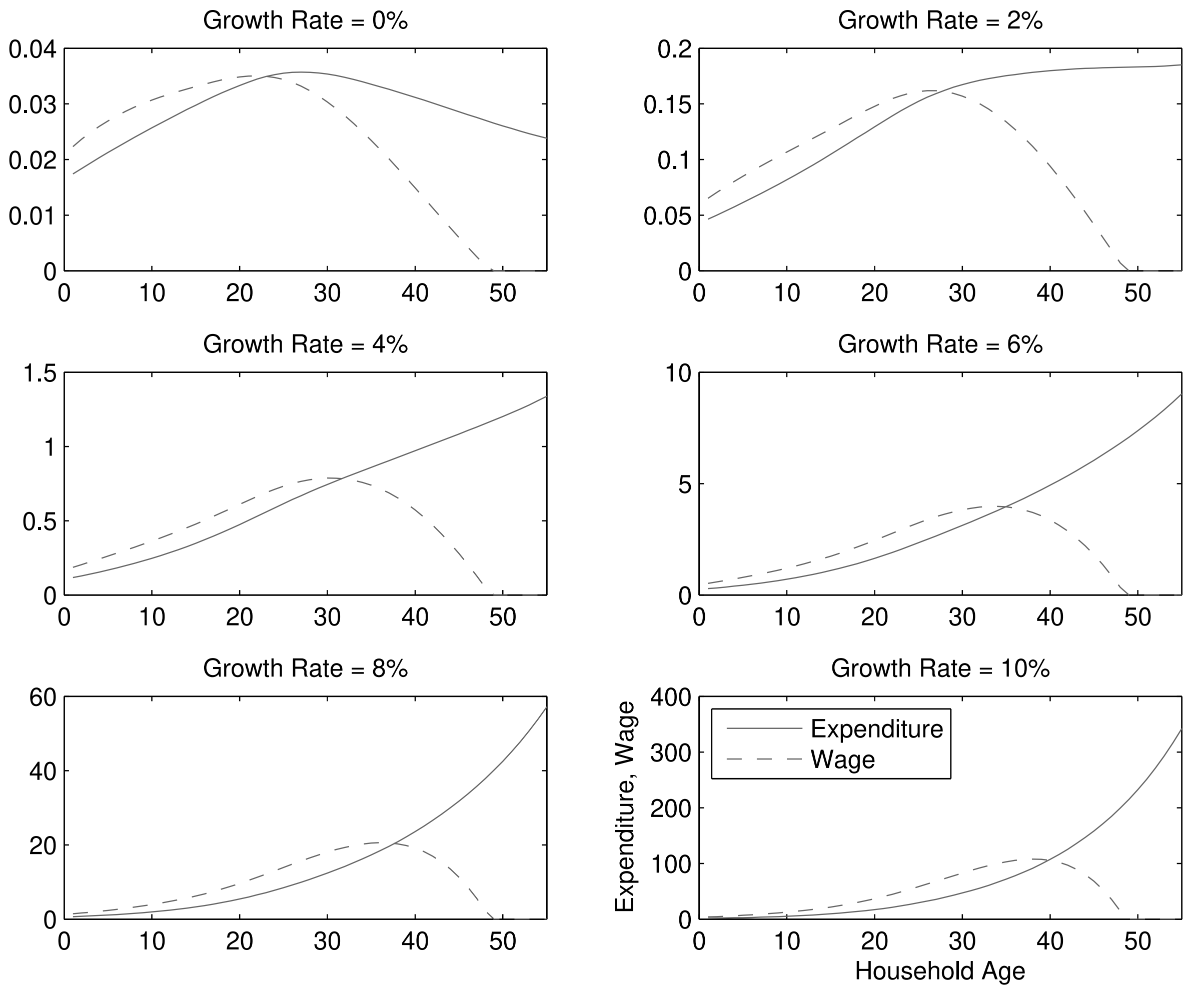

Figure 12: Household life cycle consumption, expenditure and wage, with relative consumption ($\gamma = 1$)

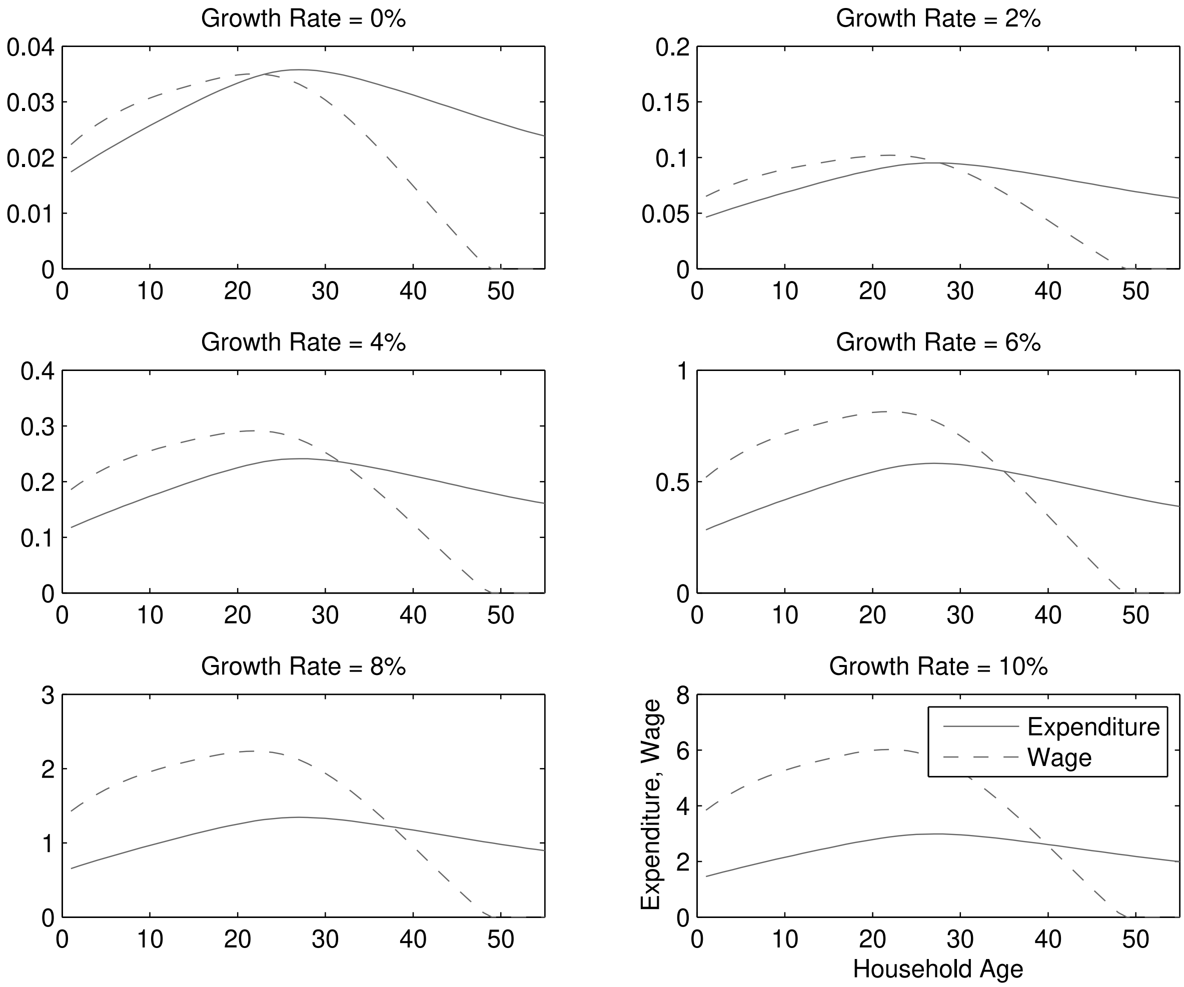

Figure 13: Cross section of household expenditure and wage, with relative consumption ($\gamma = 1$)

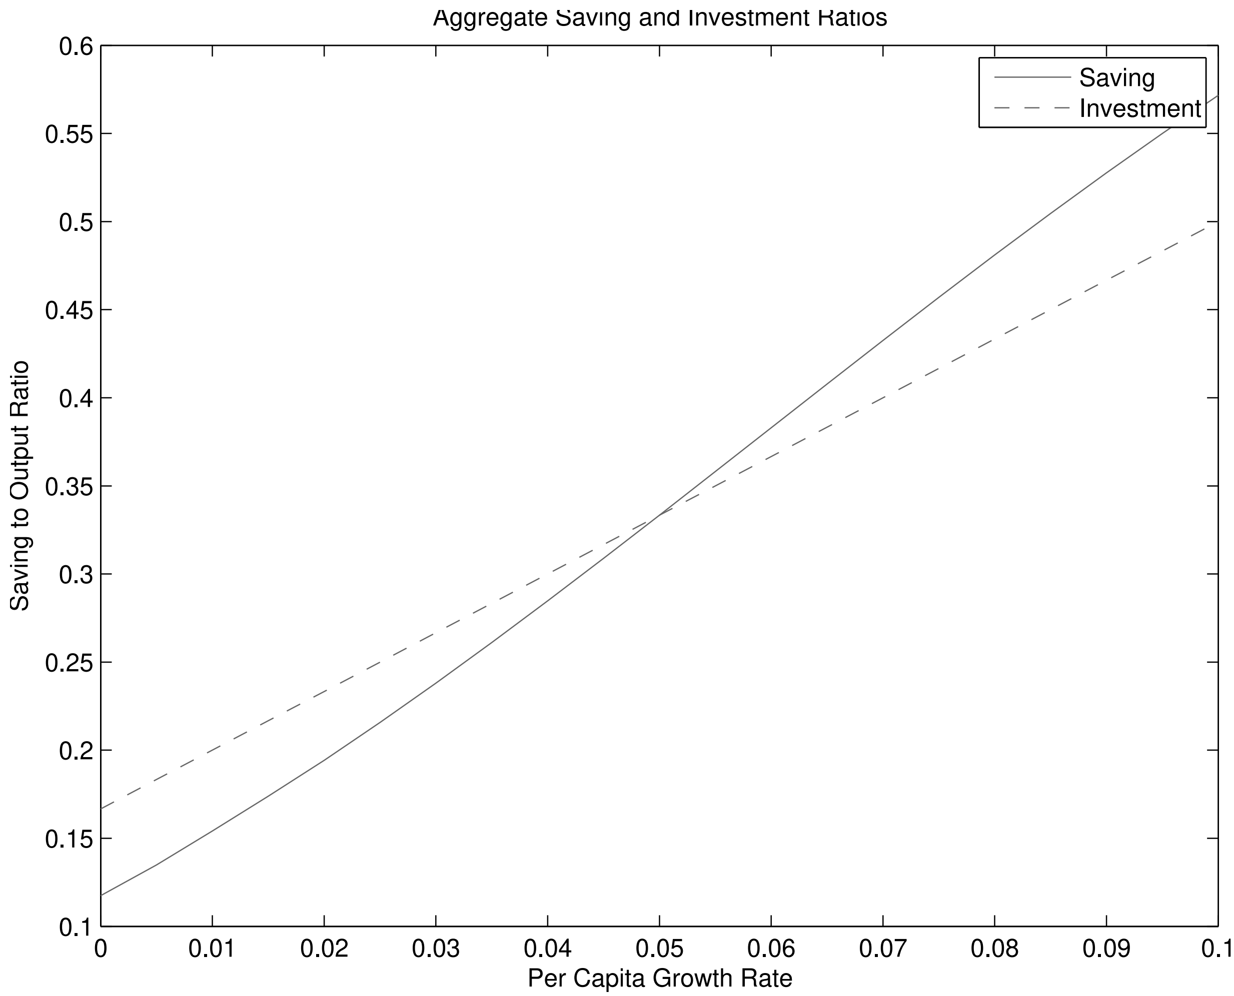

Figure 14: Saving and growth with relative consumption ($\gamma = 1$)

to simulate well the three stylized facts stated in the introduction: high growth economies enjoy high rates of saving, the shape of the consumption-age cross section is unaltered, and life cycle expenditure profiles increase more with age for higher growth rates. In addition, the simulation demonstrates that with utility from relative consumption high growth economies should hold trade surpluses, as is often the case empirically.

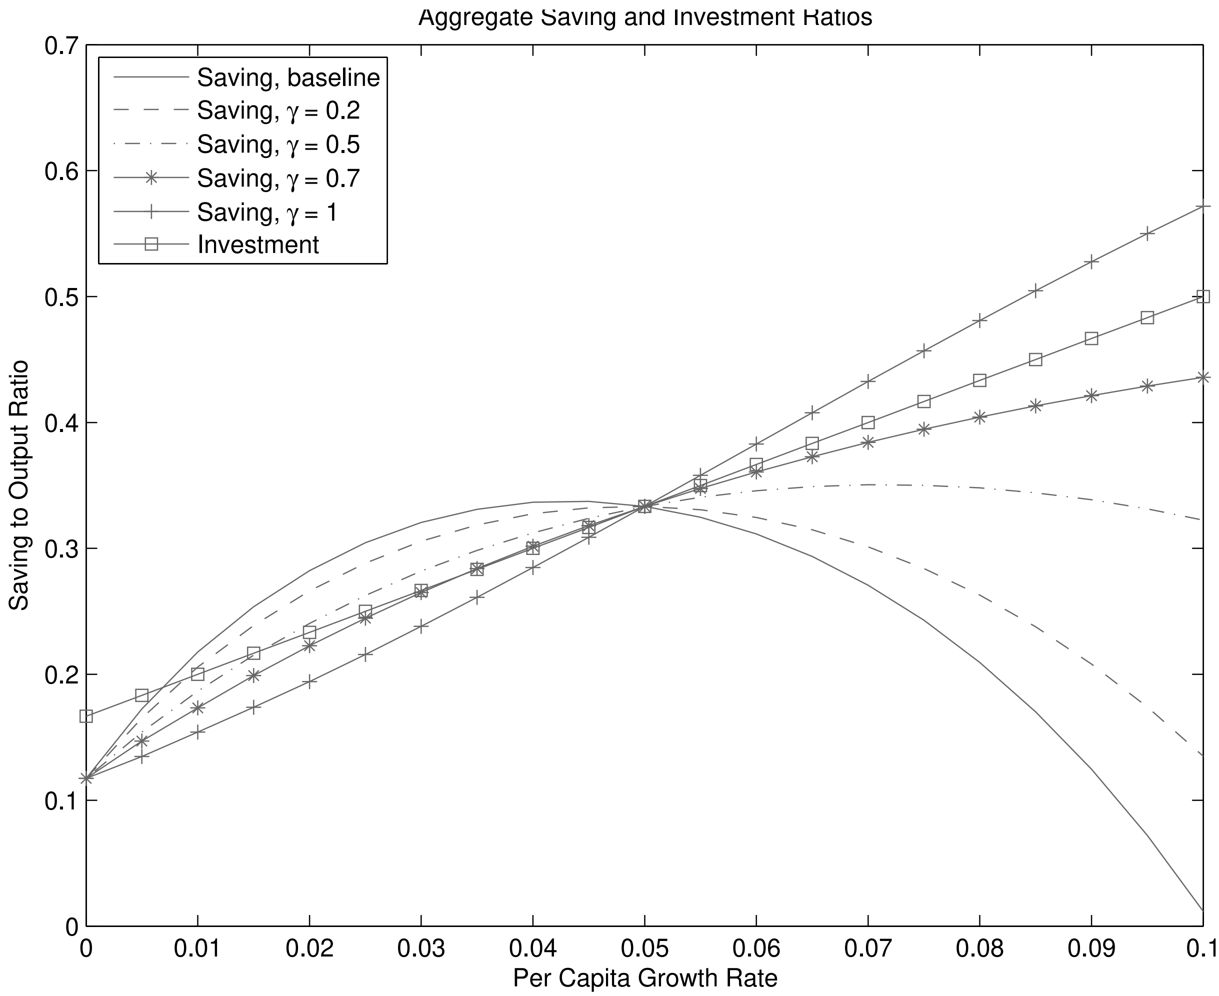

Figure 15: Saving and growth - a sensitivity analysis for different levels of relative consumption (γ)

Figure 15 presents a sensitivity analysis for different levels of γ, from zero (the baseline case) to one. The figure clearly indicates that to get a strong positive relation between economic growth and the saving rate through the relative consumption channel γ needs to be close to one. Results presented in Gruber (2018) suggest this is indeed the case, as it shows relative income, rather than absolute income, to be the determinant of household saving behavior.

6 Conclusion

High household saving rates in rapidly growing developing countries have been key factors in determining international capital and trade flows for the several past decades. The excess saving created by growth allowed poorer but rapidly developing countries to maintain a trade surplus and 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 as well.

Yet mainstream utility functions – which play a key part in current macroeconomic modeling – not only do not account for this phenomenon but in fact predict the opposite: high growth countries are supposed to have low saving rates, and effectively borrow from the rest of the world.

This paper examines the ability of a standard life-cycle model to replicate the three central stylized facts of the growth-saving relation: a positive relation from aggregate growth to the aggregate household saving ratio (with growth fully affecting the income path of individual cohorts); a correlation between the aggregate growth rate and the household expenditure growth rate over the life cycle; and a consumption-age cross-sectional distribution that is largely invariant with regard to aggregate growth.

As demonstrated, even when adding the elements of growth uncertainty and credit constraints, the life-cycle model is unable to match the three stylized facts above. However, utility from relative consumption presents itself as a viable explanation for all three. These results highlight the inadequacies of current mainstream macroeconomic modeling frameworks in dealing with high-growth environments.

While this paper focuses on utility from relative consumption as a key empirically-supported element that is able to replicate economic behavior in high-growth environments, other explanations cannot be ruled out. It is important that the macroeconomic research community become more aware of the the serious shortcomings of its current tool box, as exposed in the rapid growth environment, and set out of rectify them.

References

Acemoglu, Daron, Philippe Aghion, and Fabrizio Zilibotti, "Distance to Frontier, Selection, and Economic Growth," *Journal of the European Economic Association*, 2006, *4* (1), 37–74.

Aghion, Philippe, Diego Comin, Peter Howitt, and Isabel Tecu, "When does domestic savings matter for economic growth?," *IMF Economic Review*, 2016, *64* (3), 381–407.

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.

Brunnermeier, Markus K. and Stefan Nagel, "Do wealth fluctuations generate time-varying risk aversion? Micro-evidence from individuals asset allocation," *American Economic Review*, 2008, *98* (3), 713–736.

Carroll, Christopher D. 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.

_ **, Byung-Kun Rhee, and Changyong Rhee**, "Are There Cultural Effects on Saving? Some Cross-Sectional Evidence," *The Quarterly Journal of Economics*, 1994, *109* (3), 685–700.

_ , _ , and _ , "Does Cultural Origin Affect Saving Behavior? Evidence From Immigrants," *Economic Development and Cultural Change*, 1999, *48* (1), 33–50.

_ , 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.

Curtis, Chadwick C., Steven Lugauer, and Nelson C. Mark, "Demographic Patterns and Household Saving in China," *American Economic Journal: Macroeconomics*, 2015, *7* (2), 58–94.

Deaton, Angus 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.

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.

Fernández-Villaverde, Jesús and Dirk Krueger, "Consumption over the life cycle: Facts from consumer expenditure survey data," *The Review of Economics and Statistics*, 2007, *89* (3), 552–565.

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.

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

Gruber, Noam, "Keeping Up with the Zhangs: Relative Wealth and Household Saving Behavior," *Journal of Macroeconomics*, 2018, *55*, 77–95.

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

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

Jin, Ye, Hongbin Li, and Binzhen Wu, "Income Inequality, Consumption and Social-Status Seeking," *Journal of Comparative Economics*, 2011, *39*, 191–204.

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.

Lee, Ronald, Sang-Hyop Lee, and Andrew Mason, "Charting the Economic Life Cycle," *Population and development review*, 2014, *34* (Suppl 1), 208.

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.

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

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.