Precautionary Saving of Chinese and US Households

Horag Choi[a] Steven Lugauer[b*] Nelson C. Mark[c]

July 2016

Abstract

We employ a model of precautionary saving to study why household saving rates are so high in China and so low in the US. The use of recursive preferences gives a convenient decomposition of saving into precautionary and non precautionary components. This decomposition indicates that over 80 percent of China's saving rate and nearly all of the US saving arises from the precautionary motive. The difference between the US and Chinese household income growth rates is vastly more important than differences in income risk for explaining the saving rates. The key mechanism in the model is that precautionary savers have a target wealth-to-income ratio and rapid income growth necessitates a high saving rate to maintain the ratio.

Keywords: Precautionary Saving, Recursive Preferences, China

JEL: E2, J1

[a] Department of Economics, Monash University

[b] Department of Economics, University of Kentucky

[c] Department of Economics, University of Notre Dame and NBER

*Corresponding author. University of Kentucky, Department of Economics, Gatton School of Business and Economics, Lexington, KY, 40506, USA. stevenlugauer@gmail.com. (574) 485-8535.

1 Introduction

From 1989 to 2009, the household saving rate in China averaged slightly over 20 percent of disposable income. This saving rate is high by international standards and contrasts sharply with the 4 percent average in the US over roughly the same time period. Both policy makers and academics have developed an interest in understanding the divergent saving rates in the world's two largest economies. In China, the high saving rate helps to drive investment-led growth, contributes to its external imbalance, and presents a challenge to its goal of 'rebalancing' growth towards consumption. The low saving rate in the US, on the other hand, presents its own set of challenges for adequate provision of old-age support and a deepening of international indebtedness. What accounts for such a large discrepancy in saving rates between the two countries, and can a single model provide an explanation? In this paper, we study the extent to which a model of precautionary saving can answer these questions.

In the policy domain, Bernanke (2005) identifies the precautionary motive as the underlying source of high Asian saving rates and for generating external imbalances that depressed the world interest rate. This so-called 'savings glut' also may have contributed to the sub-prime crisis. When opportunities for unsecured borrowing and risk-sharing are limited, the story goes, households use asset accumulation to serve as a precautionary buffer against zero or very low income states of nature. We show that both income growth and income risk at the household level are much higher in China than in the US. Since the observed differences in the properties of household income for the two countries are stark and conspicuous, it seems reasonable to ask whether these differences in income characteristics can generate the divergent saving rates observed in the data.

Our structural model of saving decisions endows households with recursive Epstein-Zin-Weil preferences (Epstein and Zin 1989, Weil 1989). These preferences generalize power utility by treating the intertemporal elasticity of substitution (IES) and the coefficient of relative risk aversion (RRA) as separate parameters. Following the precautionary saving literature (originally developed for and applied to the US), households in the model hold a positive 'buffer' of assets and face idiosyncratic income risk.

We find that riskier income does induce higher saving rates in the model, which is no surprise. What is surprising is that the differential income risk between China and the US explains a relatively small portion of the Chinese-US saving rate gap. It turns out that the difference in the household income growth rates across countries is a much more important factor. We show this relationship between income growth and saving by embedding estimates of US and Chinese income dynamics into simulations of our model of household saving.

Our analysis of the data and model yields a number of insights. To help organize the discussion, we refer to the following list of six results, with number 6 as the main finding.

1. Chinese households save a far higher share of their income than US households.
2. Household income in China grows at a faster rate than in the US.
3. Chinese households face more severe income shocks than US households.
4. A large portion of saving in both China and the US is precautionary saving (in the model).
5. Income growth and income risk (and IES and RRA parameters) affect the amount of precau-

tionary saving.

6. The difference in income growth rates between China and the US is quantitatively more important than the difference in income risk as an explanation for the US-China gap in household saving rates.

Result 1 is widely known, and explaining this big difference in aggregate household saving rates (20.5 percent versus 4 percent) motivates our analysis. We arrive at results 2 and 3 by estimating income growth, transitory and permanent income shocks, and the probability of a large negative income shock separately for each country following the procedure laid forth in Carroll (1992). The data used to estimate the income dynamics for Chinese households comes from the *China Health and Nutrition Survey*. The analogous estimation for US households employs the *Panel Study for Income Dynamics*. The findings are striking; Chinese households enjoy much higher income growth rates on average (7.3 percent versus 0.6 percent), yet face more severe income shocks along each dimension.

We use our model of household saving decisions to assess the role played by income dynamics (results 2 and 3) in explaining saving behavior (result 1). The online appendix (Choi et al. 2016) characterizes the relationships analytically. Complete closed form solutions are not available, but the model equations suggest that additional income risk increases precautionary saving. The effect of income growth on saving is not necessarily monotone, and the IES and RRA preference parameters also have an ambiguous impact on saving. Therefore, we simulate the model economy for a wide range of parameter values, taking the income dynamics as exogenous.

The model's recursive utility structure provides a convenient way to quantify the size of the precautionary component of the saving rate. The non-precautionary part of saving equals the implied saving rate when people are risk-neutral. The gap between the actual and risk-neutral saving rate is then due to the precautionary motive. By this accounting, the risk-neutral saving rate in the US is nearly zero. Almost all of the US saving rate can be attributed to the precautionary motive in our model. Risk-neutral saving by Chinese households is also small relative to their large total saving. Thus, result 4 is that most of the saving within the model is precautionary.

We also use the simulations to calculate how saving rates change in response to varying the parameters governing the income process and preferences. Result 5 is, in short, that these parameters matter. Further, the precautionary saving model can replicate the observed aggregate Chinese and US household saving rates using several configurations of the IES (σ) and RRA (γ) parameters. Thus, to further discipline the model, we estimate the preference parameters by the method of simulated moments using cross-sectional consumption data from the *Urban Household Survey* for China.

For China, we estimate (γ, σ) to be (3.8, 2.2). For the US, we set the preference parameters (γ, σ) to (7.5, 1.5), estimates found in the literature (see Bansal and Yaron 2004).[1] At these parameter values, the model predicts a saving rate of 26.5 percent for China and 3.6 percent for the US. Again, most of the saving is precautionary in nature for both countries, at these specific parameter values. Moreover, the large US-China saving rate gap is driven primarily by the difference in income growth rates. For example, if we re-simulate the model with Chinese households receiving only the low, US rate of income

[1] In Choi et al. (2014), we also estimate the US parameters via the method of simulated moments, obtaining values close to Bansal and Yaron (2004).

growth (and keep all other parameters fixed at their Chinese values), then the Chinese saving rate falls to 4.3 percent (near the US rate). To reiterate, the difference in income growth rates (0.6 versus 7.3 percent) can account for nearly all the difference in saving within the model. This finding is result 6, our main finding.

Result 6 is surprising, and the intuition behind it requires illumination. Precautionary savers, as discussed by Carroll (1992, 1997), have a target asset to income (or consumption) ratio. If this target ratio is relatively insensitive to the rate of income growth, the saving rate in a high-growth country must be high in order for assets, in the numerator, to grow at the same rate as income, in the denominator. This mechanism connecting household income growth to saving accounts for a large part of the difference in saving rates across the two countries.

Note, though, the relationship between income growth and saving is not always monotonic. Our analysis of the model equations show (and the model simulations confirm) that higher income growth can potentially decrease household saving, depending on the IES and RRA parameter values. Two additional channels interact to push down saving in response to income growth. First, agents expecting higher income in the future want to save less because of the consumption smoothing motive. Second, income growth makes households less vulnerable to some types of income risk leading them to lower their target wealth-to-income ratio. In our simulations at the estimated parameter values, however, the accumulation of assets for precautionary saving offsets the two effects in the other direction, especially for China.

The model predicts that China's household saving rate will decrease when income growth begins to taper. On the flip side, the analysis indicates that the US saving rate could be increased substantially if household income growth could be raised - admittedly, not an easy task. Our findings also help explain why saving rates are often high in high growth countries, as observed in Japan during the 1960s-1970s, Korea during the 1980s-1990s, and currently in China. The positive empirical relation between growth and saving has posed a puzzle in the sense that consumption smoothing arguments lead us to expect higher future income growth to depress current saving (Carroll et al. 2000).

The high Chinese household saving rate has generated an active area of research. As saving is a multi-faceted phenomenon, research has also formed a multi-faceted approach. Mechanisms studied include capital market imperfections in China (Horioka and Terada-Hagiwara 2012 and Coeurdacier et al. 2012), rapidly increasing housing prices (Wan 2015 and Wang and Wen 2012), and a speculative bubble more generally (Wan 2016). Kraay (2000) and Meng (2003) incorporate precautionary saving aspects into their empirical analyses and Chamon and Prasad (2010) estimate the effects of the shifting state-to-private burden of educational and medical expenditures on the urban saving rate. More closely related to our paper is Chamon et al. (2013). They show that the observed increase in the volatility of transitory income shocks (from downsizing of the State-Owned sector) has led to over a 5 percentage point increase in the saving rate. He et al. (2014) use China's large-scale reform of state-owned enterprises in the late 1990s as a natural experiment to identify exogenous variation in income uncertainty and to estimate the importance of precautionary saving in China. Their estimation suggests that precautionary saving accounts for about 30 percent of the wealth accumulation by urban households. We also find that precautionary saving is quantitatively important, but using a different methodology. Thus, our results are similar to He et al. (2014) in this regard.

We do not attempt to incorporate all these theories explicitly into a single model; rather, our estimates capture the income process faced by households in a reduced form way, without taking a particular stance on how the income uncertainty and income growth arise. Other research has examined the relationship between income growth and saving rates. Carroll et al. (2000) show that saving rates can increase with income growth when households have preferences with habit persistence, so when income growth accelerates, households that are accustomed to a low level of consumption increase their saving. Chen et al. (2006) build a general equilibrium model in which the rate of technology growth affects saving via the interest rate. In contrast, we show that increased growth induces higher saving rates for a given interest rate. Song and Yang (2010) focus on saving and income by age in China, in the cross section. Additionally, several studies have focused on the role of life-cycle effects and demographic variation on the saving rate. Modigliani and Cao (2004), Horioka and Wan (2007), and Horioka (2010) undertake empirical analyses that test for the significance of these channels; whereas, Banerjee et al. (2010), Curtis et al. (2015), Song et al. (2015), and Choukhame et al. (2014) quantitatively model how demographic changes affect China's saving rate through life-cycle channels. Wei and Zhang (2011) argue that the high saving can be partly explained by the sex-ratio imbalance that has emerged as an unintended consequence of the population control measures.

These competing explanations are complementary rather than exclusive. Each makes some contribution to explaining high Chinese saving rates, but it is unlikely that any single mechanism is the sole explanation. In our study of the importance of the precautionary saving motive, one might ask if our mechanism is robust to alternative explanations proposed in the literature. While it is not feasible (nor necessarily desirable) to build a model that incorporates all facets of saving, we can establish robustness of the interaction between precautionary saving and income growth to demographic, family size, and life-cycle considerations. To do this, we modify the baseline model to include both the demographic structure and age-dependent heterogeneity in income risk. To impose the observed age distribution for each country, we give the model agents finite life-times and calculate their saving at different ages. Agents spend their last periods of life in retirement living off of accumulated assets and a pension. We also modify the utility function to incorporate explicit care of dependent children (following Curtis et al. 2015). As in the previous studies mentioned above, the age distribution has a large effect on saving. However, the precautionary component to saving remains large, too. Both channels contribute to China's high household saving rate.

While work that focuses directly on precautionary saving in China is relatively thin, it has been and continues to be a richly studied topic for US households. Gourinchas and Parker (2002) combine life-cycle and precautionary motives in a quantitative model, while Parker and Preston (2005) use cross-sectional data to estimate an empirical decomposition of the saving rate into precautionary and non-precautionary components. The general methodology of our paper aligns with the analyses of Zeldes (1989), Deaton (1991), and Carroll (1992, 1997). These studies work with constant relative-risk aversion utility. Therefore, an additional contribution of our analysis comes from using the more general

Epstein-Zin-Weil preferences.

The remainder of the paper is laid out as follows. The next section discusses the aggregate household saving rates for China and the US, which we take as target values for the model simulation exercises. Section 3 presents the model for infinitely-lived Epstein-Zin-Weil households who face risky income streams. Section 4 studies properties of the model through simulations. In Section 5, we estimate the preference parameters for China by the simulated method of moments in order to discipline the model. Then, we assess the ability of the model to explain the data when preference parameters are set to the point estimates. Section 6 studies the version of the model modified to include population age and family size demographics. Section 7 concludes.

2 Household Saving Rates in China and the US

Before sweeping economic reforms began in 1978, saving was relatively low in China. From 1959 to 1978, American households saved a higher fraction of income than Chinese households. This changed dramatically after 1980. Since that time, Chinese saving rates have climbed enormously while US saving rates have trended down.[2] Between 1980 and 1986, China's saving rate increased from 12 percent to 16 percent then fell to 11 percent in 1989, and has increased more or less steadily thereafter. In the US, on the other hand, the saving rate averaged 9.1 percent from 1959 to 1984, trended downward for 23 years reaching a low of 2.4 percent in 2007, before rising somewhat as people deleveraged during the great recession.

Broadly speaking, China underwent two sets of economic reforms during this time period. The first-round of post-central planning reforms began in 1978 and progressed through the 1980s. The initial reforms, directed at agriculture, led to increases in farm productivity and a surplus in farm labor. The labor released from the land fueled an explosion of private entrepreneurship formed by Township and Village Enterprises, self-employment businesses, and private-run firms often located in rural areas and leading to a rapid rise in rural income (Huang (2008), Chapter 2).

The Tiananmen Square protests in 1989 marked a critical point after which many policies were reversed and focus shifted away from the rural economy towards the urban areas. This second phase of reform centered on regulatory and administrative restructuring of key market sectors. Of special relevance to our study was the significant downsizing of the State-Owned sector, which resulted in the loss of generous health, retirement, education, and housing benefits. The transition away from state-provision of services and income insurance has been referred to as the dismantling of the "iron rice bowl", and the changes have created new incentives for precautionary saving by households (Chamon and Prasad (2010), He et al. (2014)). This phase of the reform process is still ongoing.

The Chinese household survey panel data that we use to estimate the income dynamics are available from 1989 to 2009. Hence, we focus on the saving rate over this span of time, the period covered by the second set of Chinese economic reforms. To facilitate a comparison over (approximately) the same time

[2] Personal saving rate as a percentage of disposable income. Sources: US Department of Commerce, Bureau of Economic Analysis and various issues of the *China Statistical Yearbook*.

Table 1: Aggregate Household Saving over the Time-Span of Interest

	CHINA	USA
Average Saving Rate	20.5	4.0

period, we select waves of panel data from 1992 to 2007 for US households. During these periods, as a percentage of income, the average aggregate Chinese household saving rate was 20.5 percent whereas the average US rate was a much lower 4.0 percent. Table 1 summarizes result 1.

Result 1: Chinese households save a far higher share of their income than US households.

The question we ask of the model is whether it can explain these average saving rates. Our focus is not on the evolution of saving rates over time, as the model is not equipped for transition dynamics.[3] However, the potential factors (e.g. income growth and risk) that we consider also may have contributed to the observed patterns in the time series. For example, both household income volatility and wage growth have increased along with China's saving rate.[4] In our model simulations, the difference in saving rates between China and the US is primarily driven by China's high income growth rate. Lower growth generally (but not always, depending on the exact parameters used) leads to less saving in the model, which is consistent with the time series facts. Next, we present the model.

3 A Model of Household Saving

We begin with a description of the household's exogenous income process. We estimate the income process using micro data. The large disparity in the estimates of the income growth rates and income shocks across the two countries comprise results 2 and 3. Subsection 3.2 presents household preferences. The online appendix (Choi et al. 2016) explores the household saving decisions analytically within the model and shows how the model is transformed to induce stationarity for the simulation exercises.

3.1 The Income Process

We use the permanent-transitory income component model of Zeldes (1989) and Carroll (1992, 1997). Households draw different realizations of exogenous labor income from the same initial distribution. The expected income growth rate is common across individuals within a country. Households are infinitely lived and experience idiosyncratic realizations of permanent and transitory income shocks in each period t.[5]

[3] For analyses that do focus on the evolution of the saving rate in China, see Curtis et al. (2016), Coeurdacier et al. (2012), and Song and Yang (2010).

[4] Note, the pattern of household income and wage growth has been quite distinct from GDP growth in China. In recent years, wage growth has generally increased while output growth has slowed. See Yang et al. (2010) for more on wages within China. Similarly, Chamon et al. (2013) show that the variance of income has been increasing in China.

[5] See Aiyagari (1994) for a related model with only transitory income.

Markets are implicitly incomplete in the model; agents cannot borrow, nor can they purchase contingent claims or other insurance instruments to diversify away from the labor income risk. In reality, insurance markets and other social safety nets are well developed in the US. Therefore, what we seek to measure in the data below (for both the US and China) are shocks to income net of all such transfers (including those between family members). The household income risk faced by the model agents should be interpreted as driven by those remaining shocks that cannot be diversified away.[6]

Let $Y_{i,t}$ be "household income" for individual $i \in [1, N]$. Income is comprised of a permanent part $P_{i,t}$ and a transitory part $e^u_{i,t}$ and evolves according to

$$Y_{i,t} = P_{i,t} e^{u_{i,t}}. \tag{1}$$

Log permanent income $\ln(P_{i,t})$ follows a random walk with drift. In levels, it evolves according to

$$P_{i,t} = e^g P_{i,t-1} e^{n_{i,t}}, \tag{2}$$

where g is the common growth rate of household income.[7] The term $e^{n_{i,t}}$ represents the log normally distributed innovation to permanent income where $n_{i,t} \overset{iid}{\sim} N(\mu_n, \sigma_n^2)$ with $\mu_n = -\sigma_n^2/2$.

The transitory component is drawn from a mixture of a lognormal variate and a low probability event, which translates into zero income for that year. It evolves according to

$$u_{i,t} = \begin{cases} N\left(\mu_u, \sigma_u^2\right) & \text{with probability } (1-p) \\ -\infty & \text{with probability } p \end{cases},$$

where p is the probability of drawing zero income and $\mu_u = -\sigma_u^2/2 - \ln(1-p)$.[8] This mixture of distributions is often employed in analyses of precautionary saving because the income data appears to be distributed log normally except for a concentration of observations at the lower tail of the income distribution.[9] The permanent and transitory shocks are assumed to be orthogonal to each other, $Cov(u_{i,t} n_{i,t}) = 0$. We estimate the four parameters $(g,\ \sigma_n,\ \sigma_u,\ p)$ governing the income process from household-level data separately for the US and China.

Household Survey Data. The Chinese data comes from the *China Health and Nutrition Survey* (CHNS), which contains income information for a panel of households in the years 1989, 1991, 1993, 1997, 2000, 2004, 2006, and 2009. The survey relies on a multistage random cluster process to track about 4,400 households, varying in terms of geography and socioeconomic status.[10] We set "labor

[6] This interpretation of income volatility is used in most of the precautionary saving literature. See Carroll (1992, 1997). Zeldes (1989) makes similar arguments concerning borrowing constraints.

[7] Following the literature, we keep g fixed. A key for our results is that g has been much lower in the US than China over an extended period of time and that households expect income growth to stay more or less the same in the near future. We think this is a reasonable set of assumptions on g.

[8] We follow Carroll (1997) by setting the mean of $n_{i,t}$ to $-\sigma_n^2/2$ so that $E\left(e^{n_{i,t}}\right) = e^{\mu_n + \frac{\sigma_n^2}{2}} = 1$ and by adjusting the mean of the transitory shock so that $E\left(e^{u_{i,t}}\right) = 1$.

[9] Carroll (1992) contains a more complete explanation for this income process. Both the US and Chinese data we employ appear to be distributed log-normally except at the lower end of the distribution.

[10] Detailed information on the survey can be found at www.cpc.unc.edu/projects/china. The survey contains information at the individual and household level. We aggregate to the household level for our analysis.

income" equal to total household non-capital income, including income earned by any family member and any transfer payments.[11] This measure of income most closely resembles the concept of income in the model developed below. We restrict observations to households in which the same individual was the head of the household for each year (for which data exists) and for which the head was older than 24 and younger than 60, with complete data on education and occupation.

The US data comes from the *Panel Study of Income Dynamics* (PSID). We impose the same data restrictions used for the Chinese data. To make the time-spans covered for US households roughly comparable to the Chinese data, we use data from the 1992, 1994, 1998, 2001, 2005 and 2007 waves of the PSID.

Estimates of the Income Process Parameters. Estimation of the income process parameters (g, σ_n, σ_u, p) follows Carroll (1992). The strategy is to remove aggregate time trends, predictable life cycle or occupation dependent fluctuations, and household fixed effects (hence the need for panel data) from the income data, then use the remaining variation to estimate the parameters (σ_n, σ_u, p) characterizing the shocks to the income process. The growth rate (g) of household income is calculated as the average real growth of income across households over the entire sample period (see the online appendix for more detail). Table 2 reports parameter estimates for the income process and compares our estimates to those reported by Carroll (1992) for an earlier time period (1968-1985) in the US and by Chamon et al. (2013) for China.[12]

Our estimated average income growth g for the Chinese households is 7.3 percent per year. China's income growth is high and evidently quite risky. The estimated probability of suffering a near zero income event p is 2.24 percent.[13] The estimated standard deviation for shocks to the transitory and permanent components of income are $\sigma_u = 0.580$ and $\sigma_n = 0.127$. Our estimates of the income shock process are similar to those reported in Chamon et al. (2013) based on the same data, but a different estimation methodology.

In sharp contrast, households in the US PSID sample experienced very low real income growth but were exposed to much less risk. Over our sample, the estimated growth rate g is just 0.60 percent per year from 1992 to 2007. The estimated probability of experiencing zero or near zero income is 0.10 percent, which is only 4 percent as large as the probability in China. These big differences encapsulate results 2 and 3.

Result 2: Household income in China grows at a faster rate than in the US.

Result 3: Chinese households face more severe income shocks than US households.

Carroll (1992) also used data from the PSID to characterize the income uncertainty of US households for the years 1968-1985 (also see Meghir and Pistaferri (2004)). Thus, we can see how the income process

[11] Total household income includes family transfers, unemployment insurance, social security, means-tested welfare, etc. The specific public policy details surrounding these programs could affect saving behavior. We abstract from such issues and focus on the properties of the realized income stream.

[12] Chamon et al. do not incorporate the near-zero income events.

[13] Following Carroll (1992), a near-zero income event is defined as annual non-capital income of less than 10 percent of the specific household's average annual income.

Table 2: Estimated Income Process

		Our Calculations		Chamon et al. (2010)	Carroll (1992)
		China	US	China	US
		1989-2009	1992-2007	2006	1968-1985
g	Income Growth (%)	7.30	0.60	n.r.	2.00
p	$\text{Prob}(e^u) = 0$ (%)	2.24	0.10	n.a.	0.65
σ_u	Transitory	0.580	0.410	0.604	0.160
σ_n	Permanent	0.127	0.121	0.121	0.126

n.r.: not reported. n.a.: not applicable.

has changed over time in the US. Compared to the 1968 - 1985 numbers reported in Carroll (1992), our results indicate that the probability of near zero income has decreased. However, our estimated transitory income shock volatility σ_u is more than twice as large as Carroll's estimate (0.41 versus 0.16). Our estimate of the standard deviation of the shock to permanent income ($\sigma_n = 0.12$) is about the same as from the earlier period. According to our estimates, about one-third of households in a given year experience a positive or negative transitory shock of at least 40 percent of their income, while a (different) third experience at least a 12 percent shock to their permanent income.

The estimated income process for Chinese households sharply differs from that for US households. Chinese households have enjoyed an income growth rate over 10 times higher than for the US (7.3 versus 0.6 percent annually). However, Chinese households face a much higher probability of a zero income shock (2.2 versus 0.1 percent), as well as larger shocks to transitory and (to a lesser extent) permanent income. The remainder of the paper focuses on quantifying how much of the large difference in household saving (result 1) can be accounted for by the differences in income growth (result 2) and income volatility (result 3) within a model of precautionary saving.

3.2 Preferences and Budget Constraints

We model household saving decisions with the Epstein-Zin (1989)-Weil (1989) recursive, non-expected utility, model of household preferences. Use of these preferences allow us to decompose saving into precautionary and non-precautionary motives, compare US and Chinese saving rates, and evaluate the relative quantitative importance of each component of the income process in generating the saving rate.

Let $C_{i,t}$ be consumption of household $i = 1, \ldots, N$ at time t. We write the utility of the infinitely-lived household as

$$V_{i,t} = \left\{ C_{i,t}^{1-\frac{1}{\sigma}} + e^{-\delta} \left[E_t \left(V_{i,t+1}^{1-\gamma} \right) \right]^{\frac{1-\frac{1}{\sigma}}{1-\gamma}} \right\}^{\frac{\sigma}{\sigma-1}}, \tag{3}$$

where $\delta > 0$ is the subjective discount rate, V is recursively defined utility and E_t is the conditional expectations operator, $\sigma > 0$ is the intertemporal elasticity of substitution (IES), and $\gamma \geq 0$ is the coefficient of relative risk aversion (RRA).

Household resources can be consumed ($C_{i,t}$) or invested in an asset ($A_{i,t}$) with gross return e^r. Households cannot borrow and face the sequence of budget constraints

$$A_{i,t+1} = (A_{i,t} + Y_{i,t} - C_{i,t})\, e^r \quad (4)$$

$$A_{i,t} \geq 0. \quad (5)$$

The household's problem is to maximize (3) subject to (4) where the exogenous "labor income" Y, is generated according to (1).

The model does not in general emit analytical solutions, so in the next section we turn to model simulations. However, in the online appendix, we derive some properties of the model. We find that a household's saving rate increases with the size of the income shocks (σ_n, σ_u, p). The impact of income growth (g) and the IES and RRA parameters (σ and γ) is non-monotonic. The online appendix also details how we transform the model to achieve stationarity.

4 Quantitative Implications

This section reports the simulated saving rates generated by embedding the income process estimates from Section 3.1 into the model presented in Section 3.2. Policy functions of the stationary model are obtained by value function iteration. The implied levels of the variables (income, assets, consumption and saving) are then obtained by "un-normalizing" the variables–that is multiplying them by permanent income. Details are contained in the online appendix.

We first study the model's properties using parameter values with ranges typically assumed or estimated in the literature. In doing so, we are able to find admissible parameter values under which the model households save at rates similar to those observed in the data.

4.1 Preference Parameter Values

The parameters for the income process (g, σ_n, σ_u, p) come from the estimates based on the US and Chinese household survey data as reported in Table 2. To simulate the model, we also need values for the three preference parameters (σ, γ, δ) and the interest rate r. There exists a large literature aimed at estimating relative risk aversion and intertemporal elasticity of substitution but without much agreement and mostly using CRRA utility. The choices for σ and γ draw on estimates reported in the literature.

We consider values of the risk aversion coefficient γ between 0 and 8 to be admissible. Studies using survey data generally find γ to be in this range ($0 < \gamma < 10$ in Dohmen et al. (2005), $7.18 < \gamma < 8.59$ in Eisenhauer and Ventura (2010), and $5 < \gamma < 10$ in Vissing-Jøgensen and Attansio (2003)). Studies that use asset pricing data and constant relative risk aversion utility (which cannot separate risk aversion from intertemporal substitution) typically obtain vastly larger values, which we ignore.

Studies that estimate the intertemporal elasticity of substitution typically report values between 0.2 and 1. Table 3 lists a collection of IES estimates reported in the literature. A restrictive feature of most existing studies, however, is the assumption of CRRA utility. Among these, the Beaudry and Van Wincoop (1996) and Gruber (2006) studies obtain relatively high estimates of the IES. Recent studies by Bansal and Yaron (2004), Chen et al. (2007), Colacito and Croce (2011), and Bansal and Shaliastovich (2013) employ recursive preferences and all feature values greater than 1. On the basis of these empirical studies, we consider three values for the $IES = (0.5, 1.5, 2.0)$.

Table 3: Estimates of the Intertemporal Elasticity of Substitution from the Literature

Authors	Special Features	Range of Estimates
Biederman and Goenner (2007)	30 Year Investment Horizon	$(0.2, 0.8)$
Felices (2005)	British Household Panel Survey	$(0.05, 0.17)$
Hall (1988)	Aggregate US Consumption	0.2
Noda and Sugiyama (2010)	Japanese Data	$(0.2, 0.5)$
Ogaki and Reinhart (1998b)	Long-run Data	$(0.27, 0.77)$
Ogaki and Reinhart (1998a)	Durable Goods	$(0.32, 0.45)$
Patterson and Pesaran (1992)	UK and US Data	0.3
Skinner (1985)		$(0.2, 0.5)$
Vissing-Jørgensen (2002)	Holders of Stocks, Bonds, and No Assets	stocks: $(0.3, 0.4)$ bonds: $(0.8, 1)$ no assets: 0.2
Beaudry and Van Wincoop (1994)	State-Level Consumption	1
Gruber (2005)		2
Chen et al. (2007)	Recursive Preferences	$(1.11, 2.22)$
Bansal and Yaron (2004)	Recursive Preferences	1.5
Colacito and Croce (2011)	Recursive Preferences	2
Bansal-Shaliastovich (2013)	Recursive Preferences	1.81

We take the real interest rate for China to be 1.6 percent per annum. This figure is the average real interest rate on bank deposits from 2003 to 2012 provided by the World Bank's *World Development Indicators*. For the US, we set the real interest rate at 1.19 percent which we obtain from the average real three-month treasury bill rate from 1992 to 2007. We use the standard rate of time preference $\delta = 0.0417$ for the US, and in anticipation of our own parameter estimate, we set $\delta = 0.022$ for China, situating households to be relatively impatient.[14]

4.2 Simulated Saving Rates

For each experiment, we simulate income and saving decisions for $N = 20,000$ households. Household saving equals the aggregate of total income (summed over the 20,000 households) minus aggregate con-

[14] High impatience has often been assumed in models of precautionary saving. We will actually generate an estimate for China below, but for now we impose the values for δ.

sumption, where total income equals the sum of labor income and interest income on assets. Households begin with zero wealth and build up their target wealth-to-income ratios over time. The saving rate is initially high as households accumulate towards their target asset-to-income ratio. For China, saving and asset ratios stabilize relatively quickly after about 20 periods. For the US, stabilization of the ratios requires 50 to 60 periods. The ratios we report are stabilized values and reflect the steady-state saving rates and asset ratios.[15] Compared to China, households in the US have an income process with a lower chance of receiving zero income, a slightly smaller variance in the permanent shock, a smaller variance in the transitory income shock, and a much lower income growth rate.

Table 4 reports the model generated household saving rates and asset to consumption ratios. For China, the saving rate–risk-aversion profile is such that for a given value of σ, the saving rate is generally increasing in γ (the exception being at γ equal to 0). The same pattern is present in the asset-consumption ratio. For China, this ratio is monotonic over the range of γ considered. For the US, the saving rate begins to display a mild hump-shape with respect to γ at the higher RRA values. As the US has smaller volatility in income shocks, the consumption and utility volatility is relatively lower than that in China. Thus, consistent with our analytical findings in the online appendix, the saving rate can exhibit a hump shape with respect to γ. When risk aversion is high enough in China ($\gamma \geq 3$) and the precautionary motive is strong, increasing the IES (σ) leads to higher saving. For the US, increasing σ tends to decrease saving at low levels of γ.

The model can match the observed aggregate household saving rates for alternative pairs of σ and γ. For China, $\sigma = 1.5$ and $\gamma = 5$ and $\sigma = 2$ and $4 < \gamma < 5$ both produce saving rates in the neighborhood of 21 percent. For the US, $\sigma = 0.5$ or 1.5 and $\gamma \geq 5$ gives a saving rate nearing (but still a little below) 4 percent.

Many studies have used σ and γ values within the ranges that allow the model to match the saving rates. However, the results are somewhat sensitive to the parameterization, and some researchers have favored different values. For example, Hubbard et al. (1995) and De Nardi et al. (2010) consider lower values.[16] The disagreement in the literature over the IES is one reason why, in the next section, we will further discipline the analysis by estimating σ and γ using our specific model and the simulated method of moments.

The simulation results in Table 4 also imply that nearly all of the US household saving in the model is driven by the precautionary motive. When risk aversion is shut off, the implied saving rate is basically nil, in contrast to the corresponding $4.7 - 8.7$ percent rate for Chinese households.[17] With the smaller variance in transitory shocks, there is less need for US households to save for consumption smoothing purposes. The relatively low rate of income growth is also a factor. Recall that we are looking at saving rates at the steady state. US households have built up large stocks of assets, but because of the

[15] We, somewhat imprecisely, use the terms steady-state or stabilized to mean that the moments of interest for the distribution do not vary much from one period to the next, i.e. it is ergodic.

[16] These papers do not use Epstein-Zin (1989)-Weil (1989) preferences, so the comparison may not be exact.

[17] This interpretation might seem stark because the model abstracts from other saving motives. However, the comparison is informative on the relative contribution to saving of risk aversion versus smoothing and also as a way to illustrate how income growth works through risk aversion to affect saving.

Table 4: Implied Saving Rates and Target Asset Ratios

China						
	Saving Rate			Assets to consumption		
	σ			σ		
γ	0.5	1.5	2.0	0.5	1.5	2.0
0	0.047	0.056	0.087	0.656	0.683	0.651
2	0.068	0.075	0.056	0.990	0.992	0.676
3	0.077	0.094	0.107	1.128	1.348	1.714
4	0.090	0.129	0.206	1.355	2.073	4.955
5	0.112	0.228	0.326	1.783	4.668	10.046
6	0.125	0.307	0.362	1.994	7.364	12.674
7	0.153	0.345	0.367	2.582	8.907	13.260
8	0.195	0.362	0.378	3.640	9.775	14.944

USA						
	Saving Rate			Assets to consumption		
	σ			σ		
γ	0.5	1.5	2.0	0.5	1.5	2.0
0	0.005	0.003	0.003	0.978	0.432	0.382
2	0.013	0.005	0.005	2.625	0.837	0.750
3	0.028	0.009	0.008	5.809	1.664	1.356
4	0.033	0.018	0.016	6.756	3.534	3.036
5	0.035	0.027	0.024	6.909	5.134	4.613
6	0.036	0.033	0.029	6.877	6.288	5.568
7	0.037	0.036	0.034	7.011	6.799	6.467
8	0.035	0.037	0.034	6.581	7.058	6.464

low growth rate, they do not have to save much to maintain their desired wealth ratios. Note, that at $4.7 - 8.7$ percent, non-precautionary saving is still only a small part of saving for Chinese households, too. Thus, we conclude that saving (in our model) is primarily due to the precautionary saving motive, result 4.

Result 4: Most saving in both China and the US is precautionary.

The remainder of the paper considers several additional variations in the key parameters governing the model. A main take away from all the simulation exercises is that the household saving rate depends on the parameters governing both the income process and preferences.

Result 5: Income growth and income risk and the preference parameters affect precautionary saving.

Variations in Growth. While the riskiness of transitory income between China and the US differs substantially, perhaps the most dramatic contrast is in the expected growth rate of income. US households expect almost no growth in labor income, whereas Chinese households expect over 7 percent growth per year.

We examine how income growth affects saving by varying g from 0.5 percent to 7 percent while holding all other parameters of the income process at their estimated values. We consider $\sigma = 1.5$ and various values of γ. Results for $\sigma = 0.5$ and 2.0 are similar and not reported to save space. Table 5 reports the household saving rates and asset ratios for China. The saving rate displays a slight hump shaped pattern with respect to g for $\gamma = 2$, 4 or 6, but saving is generally increasing in g.

Table 6 shows the analogous results for the US. The saving rate displays a hump for $\gamma = 6$ and 8 but is non-decreasing in g for $\gamma = 0$, 2, and 4. It is possible to get Americans to save a higher fraction of income with more rapid income growth. For example, at an annual growth rate of 4 percent and $\gamma = 8$, the US household saving ratio is predicted to be 15.1 percent.

If households were to maintain a constant target wealth-to-income ratio, precautionary saving needs to rise with higher income growth.[18] But notice that the relative size of the buffer stock (assets to consumption ratio) generally declines with the growth rate, which generates a trade-off. Higher income growth makes households less vulnerable to income risk, leading them to reduce their target wealth-to-income ratio (and hence their saving rate declines). This countervailing effect can be seen in table 6, panel B; the asset-consumption ratio decreases with g. The first effect (maintaining the target wealth-to-income ratio) dominates for low income growth, and the second effect dominates only when the income growth rate is high enough. The potentially positive relationship between growth and saving is noteworthy as it provides an explanation for why high-growth economies tend to have high saving rates.[19]

[18] Using slightly different terminology, Wen (2010) makes a similar point about accumulating a buffer stock of wealth in a model of uninusurable risk and borrowing constraints. Wen argues that the liquidity premium from additional saving increases with income growth, leading to a positive correlation between growth and saving and working against the permanent income hypothesis.

[19] The relationship also might be informative for the closely related allocation puzzle outlined in Gourinchas and Jeanne (2013). Other mechanisms could also link growth to household saving, such as the habit formation highlighted in Carroll et al. (2000). An analysis of alternative models lies beyond the scope of this paper.

Table 5: Variations in China's Income Growth, Implied Saving Rates and Asset-Consumption Ratios

	γ				
g	0	2	4	6	8
			A. Saving Rate		
1.005	0.007	0.034	0.040	0.039	0.036
1.01	0.013	0.056	0.075	0.077	0.071
1.02	0.023	0.080	0.140	0.139	0.131
1.03	0.029	0.081	0.176	0.191	0.174
1.04	0.033	0.074	0.188	0.224	0.206
1.05	0.041	0.073	0.183	0.248	0.231
1.06	0.043	0.075	0.157	0.257	0.245
1.07	0.048	0.076	0.133	0.254	0.253
			B. Asset-Consumption Ratio		
1.005	1.606	8.039	9.706	10.728	11.945
1.01	1.366	7.164	9.743	11.172	12.590
1.02	1.174	5.291	9.938	11.416	13.401
1.03	0.964	3.290	8.969	11.712	13.753
1.04	0.797	2.000	7.419	11.379	13.890
1.05	0.797	1.477	5.778	10.910	13.985
1.06	0.691	1.252	3.901	9.987	13.777
1.07	0.658	1.060	2.531	8.641	13.062

Calculations assume $\sigma = 1.5$.

Table 6: Variations in US Income Growth, Implied Saving Rates and Asset-Consumption Ratios

	γ				
g	0	2	4	6	8
			A. Saving Rate		
1.005	0.002	0.004	0.017	0.026	0.031
1.01	0.004	0.007	0.027	0.049	0.059
1.02	0.008	0.013	0.034	0.080	0.099
1.03	0.010	0.016	0.035	0.093	0.151
1.04	0.012	0.018	0.035	0.093	0.151
1.05	0.013	0.021	0.036	0.082	0.158
1.06	0.015	0.023	0.037	0.072	0.159
1.07	0.016	0.025	0.038	0.067	0.146
			B. Asset-Consumption Ratio		
1.005	0.405	0.930	4.135	6.196	7.236
1.01	0.382	0.733	3.138	5.857	7.105
1.02	0.366	0.643	1.906	4.931	6.188
1.03	0.307	0.522	1.236	3.866	5.730
1.04	0.286	0.442	0.891	2.838	5.048
1.05	0.258	0.408	0.711	1.901	4.257
1.06	0.247	0.355	0.594	1.313	3.559
1.07	0.224	0.330	0.512	1.010	2.717

Calculations assume $\sigma = 1.5$.

Table 5 begins to show that the difference in income growth rates can account for a large portion of the US-China saving gap (result 6). Consider $\gamma = 6$ and $g = 1.07$ (near the estimate for China) as a specific example; these parameters generate a saving rate of 25.4 percent. If, however, g is set to 1.005 (near the estimate for the US), then the saving rate collapses to 3.9 percent. Similarly, allowing US households to receive China's income growth rate ($g = 1.07$) can increase US saving.

As already noted, though, the size (and sometimes even the direction) of growth's effect on saving depends on the IES and RRA parameters. The relationship between income growth and saving is non-monotonic. Tables 5 and 6, panel B shed further light on the hump-shaped relationship. Asset holdings relative to consumption decline when the income growth rate is increased to 7 percent for both China and the US and for all values of γ. In other words, the target wealth-to-income ratio falls, short-circuiting the precautionary reason to accumulate more wealth (i.e. to save).

Finally, note that risk-neutral ($\gamma = 0$) saving in China exceeds that of the US for any of the growth rates considered. Chinese households desire to save more (buffer stock saving) on account of the higher volatility of their income. The difference in income growth rates, however, is quantitatively more important for explaining the difference in saving.[20]

Variations in near-zero income event probability. Table 7 shows the results of varying the probability of receiving zero income p for Chinese households under alternative values of γ and $\sigma = 1.5$.[21] Increasing p can increase the saving rate; however, the effect is relatively small. Increasing p from 0.8 percent to 2.4 percent only increases the saving rate by about three percentage points for an RRA less than 6 and has little or no effect for γ equal to 6 or 8. These results are similar to those reported in Carroll (1997).[22] Since zero income shocks are rare and agents literally have forever to recover, the effect on saving is not large (although not necessarily trivial).

5 Simulated Moments Estimates of Preference Parameters

The analysis of the previous section investigated the general quantitative properties of the model by treating (δ, σ, γ) as free parameters. In this section, we estimate the preference parameters for China via the simulated method of moments and investigate the extent to which the estimated model can explain the observed saving behavior. For the US, we draw on parameter values estimated in the literature (based on alternative approaches) because US households have greater access to financial markets.

Our estimation employs only consumption data. Other research estimating the recursive utility function, e.g. Chen et al. (2007) and Bansal and Shaliastovich (2013), combine consumption and asset

[20] We have also run the model under CRRA utility, restricting $\gamma = 1/\sigma$. A risk aversion coefficient between 3 and 4 allows the model to match the US saving ratio, but this model does not come close to explaining the Chinese saving ratio for any value of γ. See Choi et al. (2014).

[21] We also have experimented with different values for the variances of the shocks to permanent and transitory income. The results are straightforward; larger variances induce higher precautionary saving. Since this channel is clear and quantitatively not as important as income growth for determining saving rates, we do not report the full results.

[22] In Carroll (1997), raising p from 0.1 percent to 1 percent increases the saving rate by 0.6 percentage points from 0.4 percent to 1 percent.

Table 7: Variations in p and Chinese Saving Ratios

	Percent Chance of Receiving Zero Income, p				
γ	0.8	1.2	1.6	2.0	2.4
0	0.036	0.046	0.039	0.053	0.062
2	0.048	0.064	0.053	0.073	0.081
4	0.099	0.113	0.107	0.125	0.132
6	0.297	0.306	0.298	0.308	0.281
8	0.353	0.365	0.356	0.364	0.322

Calculations assume $\sigma = 1.5$.

returns data. As is well known, risk aversion must be very high to be consistent with asset returns data. Chen et al.'s estimates of the IES range between 1.11 and 2.22 while their estimates of the risk aversion coefficient are above 60. Bansal and Shaliastovich's estimates of (γ, σ) equal $(20.9, 1.81)$. Our estimation approach is consistent with our model, which only has implications for consumption moments. Therefore, we do not necessarily generate super high estimates of γ.

The data comes from a single cross section consisting of 15,039 observations from the 2007 *Urban-Rural Household Survey.*[23] The four moments used in the estimation are the mean, variance, skewness, and kurtosis of the cross-sectional distribution for household consumption. We simulate $N_s = 50,000$ individuals over many periods, t. The simulated moments are computed at $t = 20$.[24] The online appendix provides further details.

Table 8: Simulated Moment Estimates of Chinese Preference Parameters

δ (se)	σ (se)	γ (se)	Crit
0.022	2.160	3.780	1.498
(0.000)	(0.046)	(0.056)	

Table 8 shows the estimation results. Households are impatient in the sense that the rate of time preference exceeds the interest rate, $\delta > r$. The risk-aversion coefficient is not estimated to be exceedingly high (as is the case when using asset returns). The point estimate for the IES is substantially greater than 1.

Our estimates of σ make contact with the literature on long-run risk (Bansal and Yaron 2004). Models of long-run risk employ recursive preferences and can explain several asset pricing anomalies

[23] Yi Huang kindly provided the data.

[24] Recall that convergence to steady-state ratios takes 20 periods for China.

including the equity premium puzzle, the low risk-free rate, and their volatilities, provided two key ingredients are present. First, some long run risk must exist to vary the expected growth rate of consumption over time. Second, the size of σ matters since people must prefer early resolution of uncertainty ($\gamma\sigma > 1$). Chen et al. (2007) also estimate this elasticity to be greater than 1.[25] Interestingly, although we employ different data sets, our point estimates of σ are similar in magnitude to studies using asset returns and are consistent with the long-run risk framework.

For the US, we use the σ and γ estimates from Bansal and Yaron (2004), $\sigma = 1.5$ and $\gamma = 7.5$.[26] We set δ to yield an annual discount factor of 0.96.

Saving Rates and Decomposition under Estimated Parameters. Next, we ask to what extent the observed saving rates in China and the US can be explained by the model at the parameter values above, and what part of that saving rate is driven by the precautionary motive (result 4). Table 9 summarizes the findings.

Table 9: Total and Risk-Neutral Saving Rate Implied by Estimated Parameters

	Parameter Settings			Saving
Country	γ	σ	δ	Rate
China	3.78	2.16	0.022	0.265
	0	2.16	0.039	0.032
	0	1.5	0.022	0.044
US	7.5	1.5	0.042	0.036
	0	1.5	0.042	0.002

For China, the saving rate implied by the estimated parameters is 26.5 percent. The model over predicts the average saving rate during the time-span of the CHNS data but nearly matches the saving in 2007, the year of the consumption data used to estimate the preference parameters.

The model does not converge when risk-aversion is shut down holding σ and δ at the estimated values because the utility blows up. Although we are unable to obtain the exact risk-neutral ($\gamma = 0$) saving rate, we approximate that measurement from two directions. Our first approach is to raise the rate of time preference until convergence is achieved. Doing this gives a risk-neutral saving rate of 3.2 percent. Our second approximation is to lower the IES to 1.5, which gives risk-neutral saving of 4.4 percent. These experiments allow us to reasonably conclude that China's risk-neutral saving does not exceed 5.0 percent and that the precautionary component of the saving rate is over 20 percent of income in the model.

[25] However, the Chen et al. (2007) estimates of the risk-aversion coefficient are much larger and range between 17 and 60, depending on whether they use aggregate consumption or consumption of stock holders.

[26] We also conducted the simulated moments estimation for the US preference parameters based on *Consumer Expenditure Survey* data, obtaining values, $\sigma = 1.9$ and $\gamma = 9.1$, close to Bansal and Yaron (2004).

For the US, the 3.6 percent saving rate is a little below the 4.0 percent average from 1992 to 2007. Raising γ to about 9 would allow the model to match the 4 percent rate. The implied risk-neutral ($\gamma = 0$) saving rate is basically zero, suggesting that most of the household saving rate in the US is driven by the precautionary motive.[27] Hence, result 4 continues to hold.

Finally, we consider how the Chinese saving rates change as the rate of income growth is decreased, while keeping the other parameters at the estimated values for China. Table 10 lists the implied saving rates. When the rate of income growth is at the US level ($g = 1.005$), Chinese household saving is only 4.3 percent. Nearly all the difference between the aggregate saving rates in the US and China can be accounted for by the difference in income growth. The quantitatively large effect of income growth once again supports result 6, our main finding.[28]

Table 10: Variations in China's Income Growth and the Implied Saving Rates at the Estimated Parameter Values

g	Saving Rate
1.005	0.043
1.01	0.082
1.02	0.150
1.03	0.223
1.04	0.233
1.05	0.248
1.06	0.295
1.07	0.265

$\sigma = 2.16$, $\gamma = 3.78$

Result 6: The difference in income growth rates between China and the US is quantitatively more important than the difference in income risk as an explanation for the US-China gap in household saving rates.

Overall, the relationship between growth and saving is hump shaped in table 10. Saving rates begin to decrease at high levels of growth because the consumption smoothing channel inherent in the permanent income hypothesis begins to dominate.[29]

[27] Again, we are abstracting from life-cycle and other related saving motives. See the next section.

[28] Note, increasing the US income growth rate (while holding all other parameters at their US values) greatly increases the US saving rate, but not all the way to the Chinese level of saving.

[29] We also have experimented with shocking income growth and wealth holdings. See Choi, Lugauer, and Mark (2014) for these impulse reponse functions.

6 Incorporating Demographics, Life-Cycle Savings, and Pensions

The precautionary motive is just one of several reasons for saving. Is it robust to alternative mechanisms proposed in the literature? It is beyond the scope of this paper to control for every saving motive the literature has explored. However, we can show that the precautionary saving mechanism is robust to most of the demographic based explanations, including the age distribution, family size, pensions, and life-cycle effects. These are some of the most salient alternative explanations for China's high saving rate: those related to population demographics. As discussed in the introduction, several recent papers have argued that the age-distribution, working through standard life-cycle saving channels, has had a large impact on China's aggregate household saving rate over time. Changing family sizes and old-age pension support may also have affected saving behavior over time. Our paper does not consider transition paths; however, the current demographic profiles differ between China and the US, as do their pension systems. In this section, we show that our main findings are robust to the inclusion of these micro and macro demographic and life-cycle channels.

To incorporate the observed age distributions for the US and China, we make two key changes to our precautionary saving model. First, model agents have only finite life-times. Second, the utility function for working age agents explicitly includes the consumption by their dependent children. These two changes allow us to impose the observed age distributions in a simple way. We calculate the model agent's saving at different ages. Then, the aggregate household saving rate equals a weighted average of the age-specific saving rates, where the weights are determined by the actual age distribution in the cross-section for each country as calculated from United Nations population data. Adult agents choose the amount of consumption to give to their dependent children, but the age distribution is taken as given.

We write the utility of the finitely-lived household as

$$V_{i,t} = \left\{ \left(1 + \mu N_{i,t}^{\alpha+1}\right)^{\frac{1}{\sigma}} C_{i,t}^{1-\frac{1}{\sigma}} + e^{-\delta} \left[E_t \left(V_{i,t+1}^{1-\gamma} \right) \right]^{\frac{1-\frac{1}{\sigma}}{1-\gamma}} \right\}^{\frac{\sigma}{\sigma-1}}, \tag{6}$$

where C is total household consumption (including consumption by dependent children), N is the average number of dependent children per family, parameters μ and α govern the preferences over consumption by dependent children, and all other variables are defined as before. Following Curtis et al. (2015), consumption by dependent children enters parental utility via the Barro-Becker (1989) functional form. The budget constraint is

$$A_{i,t+1} = \left(A_{i,t} + Y_{i,t} - C_{i,t}\right) e^{r} \tag{7}$$

$$A_{i,t} \geq 0. \tag{8}$$

Income level Y varies by age (estimated using the same CHNS and PSID data and procedure as for the income process), creating an age-income profile. The rest of the model, including income growth and the shock process, remains unchanged except that a portion of each agents' life is spent in retirement. In the model, retirees in China receive 25 percent of their pre-retirement income as a pension; US retirees receive 75 percent. The model agents live 85 years in total. The first 20 years are spent as dependent children, consuming only what their parents provide. At age 20, children become adults and form their own household. From age 22 to 51, adults have dependent children to support. At age 63, agents enter retirement and live off their accumulated assets and pension until death.

To run the simulations, we need to select values for the Barro-Becker parameters, μ and α. We use values that map into the estimates in Manuelli and Seshadri (2009); $\mu = 0.68$ and $\alpha = -0.56$. See Curtis et al. (2015) for an extended discussion of these parameters. With this new set-up, we have re-done all the previous simulations (except the parameter estimation based on the method of simulated moments), including using the many different combinations of IES, RRA, and income growth parameters.

Overall, results 4-6 continue to hold in the finite-lived version of the model (results 1-3 are unchanged). To conserve space, we do not report the full details, but a few points are worth stressing. First, for all the relevant combinations of the parameters, the finite-lives version of the model generates higher saving rates. The higher saving occurs because both China and the US have a significant portion of their population in their saving years. The life-cycle saving motive (i.e. saving for retirement) pushes up saving rates. The effect is particularly large for China; as has been argued in the papers noted above, demographics have contributed to China's high saving rate. Related to this, the new model can match the observed saving rates at lower RRA values.

Second, while the precautionary component of saving is still large in the new version of the model, considerable life-cycle saving remains even when the RRA parameter (γ) is set to zero. Recall that in the infinite-lives version changing γ to equal zero in the US simulations reduced saving by nearly 100 percent. In the new version, eliminating the precautionary motive to save (setting $\gamma = 0$) reduces saving by 50 to 60 percent, depending on the parameter values. In other words, the saving rate is cut in half rather than falling to zero. The same pattern holds for China; a large portion of saving is precautionary (result 4), but the demographic structure effect is present, too.

Third, income growth, working through the precautionary saving motive, continues to have a large impact on household saving decisions. For example, changing the growth rate faced by Chinese households from the observed 7 percent to the US value of 0.5 percent decreases Chinese household saving by 10 to 16 percentage points for RRA values above 2. However, in the finite-lives model, a substantial gap (as much as 10 percentage points) in saving rates remains between the US and China, even in simulations in which each country has the same income growth rate. The demographics of China pushes its saving rate higher because China has a larger share of its population in its high saving years and smaller family sizes.

Finally, the impact on saving from altering the (transitory and permanent) shocks to income or the chance of obtaining zero income remains quantitatively small relative to income growth. In other words, the difference in income growth accounts for a much larger share of the US-China saving rate gap than does the difference in income volatility (result 6), as in the finite-lives model.

7 Conclusion

This paper presents six key findings:

1. Chinese households save a far higher share of their income than US households.
2. Household income in China grows faster than in the US.
3. Chinese households face larger income shocks than US households.
4. A large part of saving in both China and the US is precautionary within our model.
5. Income growth and income risk and the IES and RRA parameters affect the amount of precautionary saving.
6. The difference in income growth rates explains much of the US-China gap in saving rates.

The difference in household saving rates (result 1) is well known, and we investigate this stylized fact by embedding country-specific household income dynamics into a model of saving decisions by infinite-lived households with Epstein-Zin-Weil recursive preferences. We use survey data to estimate the income process and find that Chinese households have experienced a higher rate of growth (result 2) along with more volatility (result 3) in their income relative to households in the US. Since the social safety net in China is less comprehensive than in the US, Chinese households face substantially higher transitory income risk.

The model's recursive preference structure allows a convenient decomposition of the saving rate into precautionary and non-precautionary components. According to this decomposition, the precautionary motive drives most of the saving rate in China and nearly all the saving in the US (result 4) within our baseline model. Model simulations also demonstrate how the growth and riskiness of income, along with the preference parameters, determine the amount of precautionary saving (result 5). However, the relationship among these parameters and the household saving rate is potentially non-monotonic.

The model can generate the high current level of Chinese saving and the low level of US saving. Somewhat surprisingly, however, the higher income growth rate in China, and not the elevated income risk, accounts for most of the China-US saving rate gap (result 6). In the model, saving is increasing over a range of income growth rates as households save aggressively to maintain a desired asset-to-income ratio. This result sheds some light on the somewhat puzzling empirical fact that countries with high income growth (like China) often have high saving rates.

Our relatively simple set-up generates powerful insights into why household saving is so high in China and so low in the US. We also show that including the population age distribution and life-cycle saving does not qualitatively alter the main findings. Our analysis does abstract from a few relevant factors, and the model could be tailored further to consider other specific aspects of China's economy such as the dismantling of old-age income security and shifting medical and educational expenditures from the state to households. However, we leave this to future research.

Acknowledgments:

We thank two anonymous referees and Pok-sang Lam (the Editor) for many constructive suggestions. We thank Jim Sullivan and Yi Huang for assistance with the data and Wenqi Wu for efficient research assistance. This paper also has benefitted from presentations at the Hong Kong Institute for Monetary Research, the 2013 Midwest Economics Association Meetings (Columbus), the University of Notre Dame, the University of Nebraska at Omaha, the 2013 Society for Economic Dynamics Meetings (Seoul), and the 2014 Chinese Economics Society Australia Conference (Melbourne).

References

Aiyagari, S. Rao, 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 109, 659-684.

Bansal, Ravi and Ivan Shaliastovich, 2013. "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," *The Review of Financial Studies*, 26, 1-33.

Bansal, Ravi and Amir Yaron, 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481-1509.

Banerjee, Abhijit, Xin Meng and Nancy Qian, 2010. "The Life Cycle Model and Household Savings: Micro Evidence from Urban China," *mimeo,* MIT.

Barro, Robert J. and Gary S. Becker, 1989. "Fertility Choice in a Model of Economic Growth," *Econometrica,* 57(2), 481-501.

Beaudry, Paul and Eric van Wincoop, 1996. "The Intertemporal Elasticity of Substitution: An Exploration using a US Panel of State Data," *Economica,* 63, 495-512.

Bernanke, Ben S., 2005. "The Global Savings Glut and the U.S. Current Account Deficit," speech delivered at the Sandridge Lecture, Virginia Association of Economics, Richmond, Va., March 10.

Biederman, Daniel K. and Cullen F. Goenner, 2007. "A Life-Cycle Approach to the Intertemporal Elasticity of Substitution," *Journal of Macroeconomics,* 30, 481-498.

Carroll, Christopher D., 1992. "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 23, 61-156.

Carroll, Christopher D. 1997. "Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics,* 112, 1-55.

Carroll, Christopher D., Jody Overland and David N. Weil, 2000. "Saving and Growth with Habit Formation," *American Economic Review,* 90, 341-355.

Chamon, Marcos and Eswar S. Prasad, 2010. "Why are Saving Rates of Urban Households in China Rising?" *American Economic Journal: Macroeconomics,* 2, 91-130.

Chamon, Marcos, Kai Liu and Eswar S. Prasad, 2013. "Income uncertainty and household savings in China," *Journal of Development Economics*, 105,164-177.

Chen, Xiaohong, Jack Favilukis, and Sidney C. Ludvigson, 2007. "An Estimation of Economic Models with Recursive Preferences," *mimeo*, New York University.

Chen, Kaiji, Ayse Imrohoroglu, and Selahattin Imrohoroglu, 2006. "The Japanese Saving Rate," *American Economic Review*, 95(5), 1850-1858.

Choi, Horag, Steven Lugauer, and Nelson C. Mark, 2014. "Precautionary Saving of Chinese and US Households," NBER Working Paper No. 20527.

Choi, Horag, Steven Lugauer, and Nelson C. Mark, 2016. "Online Appendix for Precautionary Saving of Chinese and US Households," Working Paper.

Choukhame, Taha, Nicolas Coeurdacier, and Keyu Jin, 2014. "The One-Child Policy and Household Savings," *mimeo.*

Coeurdacier, Nicolas, Stéphane Guibaud and Keyu Jin, 2012. "Credit Constraints and Growth in a Global Economy," CEPR Discussion Papers No. 9109.

Colacito, Ricardo and Mariano Croce, 2011. "Risks for the Long Run and the Real Exchange Rate," *Journal of Political Economy*, 119, 153-181.

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

Curtis, Chadwick C., Steven Lugauer and Nelson C. Mark, 2016. "Demographics and Aggregate Household Saving in Japan, China, and India," NBER Working Paper No. 21555.

Deaton, Angus, 1991."Saving and Liquidity Constraints," *Econometrica*, 59,1221-48.

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

Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jurgen Schupp and Gert G. Wagner, 2005. "Individual Risk Attitudes: New Evidence from a Large, Representative, Experimentally-Validated Survey," IZA DP working paper no. 1730.

Epstein Lawrence, and Stanley Zin, 1989. "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937-969.

Eisenhauer, Joseph G. and Luigi Ventura, 2010. "Survey Measures of Risk Aversion and Prudence," *Applied Economics*, 35, 1477-1484.

Felices, Guillermo, 2005. "Intertemporal Substitution and Labour Supply in the UK: Estimates from the British Household Panel Survey (BHPS)," *mimeo,* Bank of England.

Gruber, Jonathan, 2006. "A Tax-Based Estimate of the Elasticity of Intertemporal Substitution," NBER Working Paper No. 11945.

Gourinchas, Pierre-Olivier and Olivier Jeanne, 2013. "Capital Flows to Developing Countries: The Allocation Puzzle," *The Review of Economic Studies*, 80, 1484-1515.

Gourinchas, Pierre-Olivier and Jonathan Parker, 2002. "Consumption over the Life Cycle," *Econometrica*, 70, 47-89.

Hall, Robert, 1988. "Intertemporal Substitution in Consumption," *Journal of Political Economy*, 96, 339-357.

He, Hui, Feng Huang, Zheng Liu and Dongming Zhu, 2014. "Breaking the "Iron Rice Bowl" and Precautionary Savings: Evidence from Chinese State-Owned Enterprises Reform," *mimeo*, Federal Reserve Bank of San Francisco.

Horioka, Charles Yuji, 2010. "Aging And Saving in Asia," *Pacific Economic Review*, 15, 46-55.

Horioka, Charles Yuji and Akiko Terada-Hagiwara, 2012. "The Determinants and Long-term Projections of Saving Rates in Developing Asia," *Japan and the World Economy*, 24, 128-137.

Horioka, Charles Yuji and Junmin Wan, 2007. "The Determinants of Household Saving in China: A Dynamic Panel Analysis of Provincial Data," *Journal of Money, Credit and Banking*, 39, 2077-2096.

Huang, Yasheng, 2008. *Capitalism with Chinese Characteristics: Entrepreneurship and the State*, Cambridge University Press..

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, 1995. "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 103, 360-399.

Kraay, Aart, 2000. "Household Saving in China," *World Bank Economic Review*, 14, 545-570.

Manuelli, Rodolfo and Ananth Seshadri, 2009. "Explaining International Fertility Differences," *Quarterly Journal of Economics*, 2, 771-807.

Meghir, Costas and Luigi Pistaferri, 2004. "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72, 1-32.

Meng, Xin, 2003. "Unemployment, Consumption Smoothing, and Precautionary Saving in Urban China," *Journal of Comparative Economics*, 31, 465-485.

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

Ogaki, Masao and Carmen M. Reinhart, 1998a. "Measuring Intertemporal Substitution: The Role of Durable Goods," *Journal of Political Economy*, 106, 1078-1098.

Ogaki, Masao and Carmen M. Reinhart, 1998b. "Intertemporal substitution and durable goods: long-run data," *Economics Letters,* 6, 85–90.

Noda, Akihiko and Shunsuke Sugiyama, 2010. "Measuring the Intertemporal Elasticity of Substitution for Consumption: Some Evidence from Japan," *Economics Bulletin,* 30, 524-533.

Parker, Jonathan A and Bruce Preston, 2005. "Precautionary Saving and Consumption Fluctuations," *American Economic Review,* 95, 1119-1143.

Patterson, Kerry D. and Bahram Pesaran, 1992. "The Intertemporal Elasticity of Substitution in Consumption in the United States and the United Kingdom," *The Review of Economics and Statistics*, 74, 573-584.

Skinner, Jonathan, 1985. "Variable Lifespan and the Intertemporal Elasticity of Consumption," *The Review of Economics and Statistics*, 67, 616-623.

Song, Zheng Michael and Dennis Tao Yang, 2010. "Life Cycle Earnings and Saving in a Fast-Growing Economy," *mimeo,* Chinese University of Hong Kong.

Song, Zheng Michael, Kjetil Storesletten, Yikai Wang, Fabrizio Zilibotti, 2015. "Sharing High Growth Across Generations: Pensions and Demographic Transition in China," *American Economic Journal: Macroeconomics*, 7 (2), 1-39.

Vissing-Jøgensen, Annette, 2002. "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy,* 110, 825-853.

Vissing-Jøgensen, Annette and Orazio P. Attanasio, 2003. "Stock-Market Participation, Intertemporal Substitution, and Risk-Aversion," *American Economic Review,* 93, 383-391.

Wan, Junmin, 2015. "Household Saving and Housing Prices in China," *The World Economy,* 38, 172-192.

Wan, Junmin, 2016. "The Speculative Saving Hypothesis," *mimeo,* Fukuoka University.

Wang, Xin and Yi Wen, 2012. "Housing Prices and the High Chinese Saving Rate Puzzle," *China Economic Review,* 23(2), 265-283.

Wei, Shang-Jin and Xiaobo Zhang, 2011. "The Competitive Saving Motive: Evidence from Rising Sex Ratios and Savings Rates in China," *Journal of Political Economy*, 119, 511-564.

Weil, Philippe, 1989. "The Equity Premium Puzzle and the Risk-Free Rate Puzzle," *Journal of Monetary Economics*, 24, 401-421.

Wen, Yi. 2010. "Saving and Growth under Borrowing Constraints: – Explaining the "High Saving Rate" Puzzle," *mimeo,* Federal Reserve Bank of St. Louis.

Yang, Dennis Tao, Vivian Weijia Chen, and Ryan Monarch, 2010. "Rising Wages: Has China Lost its Global Advantage?," *Pacific Economic Review,* 15, 482-506.

Zeldes, Stephen P., 1989. "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 97, 305-346.