



Income uncertainty and household savings in China[☆]



Marcos Chamon^{a,*}, Kai Liu^b, Eswar Prasad^{c,d,e}

^a Research Department, IMF, United States

^b Norwegian School of Economics, Norway

^c Cornell University, United States

^d Brookings Institution, United States

^e NBER, United States

ARTICLE INFO

Article history:

Received 9 June 2012

Received in revised form 24 July 2013

Accepted 27 July 2013

JEL classification:

D91

J3

E21

Keywords:

China

Household savings

Income uncertainty

Pension reforms

Buffer-stock savings

ABSTRACT

China's urban household saving rate has increased markedly since the mid-1990s and the age-savings profile has become U-shaped. To understand these patterns, we analyze a panel of urban Chinese households over the period 1989–2009. We document a sharp increase in income uncertainty, largely due to an increase in the variance in household income attributed to transitory idiosyncratic shocks. We then calibrate a buffer-stock savings model to obtain quantitative estimates of the impact of rising household-specific income uncertainty as well as another shock to household income—the pension reforms that were instituted in the late 1990s. Our calibrations suggest that rising income uncertainty and pension reforms lead younger and older households, respectively, to raise their saving rates significantly. These two factors account for two-thirds of the increase in China's urban household saving rate and the U-shaped age-savings profile.

© 2013 International Monetary Fund. Published by Elsevier BV. All Rights Reserved.

1. Introduction

The Chinese economy has experienced rapid growth in average household incomes but also increased uncertainty as the economy undergoes massive structural shifts. The transformation of the economy into a more market-oriented one has been accompanied by significant policy changes, including reforms to the pension system and the hardening of budget constraints on state enterprises. Our objective in this paper is to evaluate the effects of these structural and policy changes on the degree of income uncertainty at the household level and to analyze the implications for household saving rates.

This research is motivated by the significant increase in China's household saving rate over the last two decades. The average savings rate (as a share of household disposable income) for urban households

in China increased by about 5 percentage points during the 1990s and then rose sharply by another 10 percentage points over the next decade, to a level around 30% by 2010 (Fig. 1). Because urban incomes account for a substantial fraction of national income, the national household saving rate has closely tracked the urban one.¹

The increase in household savings has contributed to a rising national saving rate, which is now among the highest in the world. It is also an important determinant of China's current account surplus, which plays a key role in discussions of global current account imbalances.² The rising household saving rate at a time of high income growth is also of interest from an analytical perspective. It seems inconsistent with a life-cycle hypothesis model without strong precautionary saving motives, which would imply that future high income growth should cause households to postpone their savings. In addition to the increase in saving rates across the board, there has

[☆] We are grateful to Loren Brandt, Robert Moffitt, two anonymous referees, as well as participants at the NBER Summer Institute, China Economics Summer Institute, IMF Research Seminar, and the Workshop on China's Macroeconomy at the University of Toronto for comments and suggestions. We thank Lei (Sandy) Ye for research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect those of the institutions the authors are affiliated with.

* Corresponding author. Tel.: +1 2026235867.

E-mail addresses: mchamon@imf.org, mchamon@alum.mit.edu (M. Chamon).

¹ The household saving rate based on national income accounts data (not shown here) indicates a higher aggregate saving rate, although the trend is quite similar to that of the urban household saving rate. In most countries, there is a discrepancy between household saving rates based on survey data versus national income accounts. The reasons include differences in definitions of consumption and income, under-sampling of the rich in household surveys, etc.

² See Prasad (2011). For a complementary discussion of China's corporate savings, see Bayoumi et al. (2010).

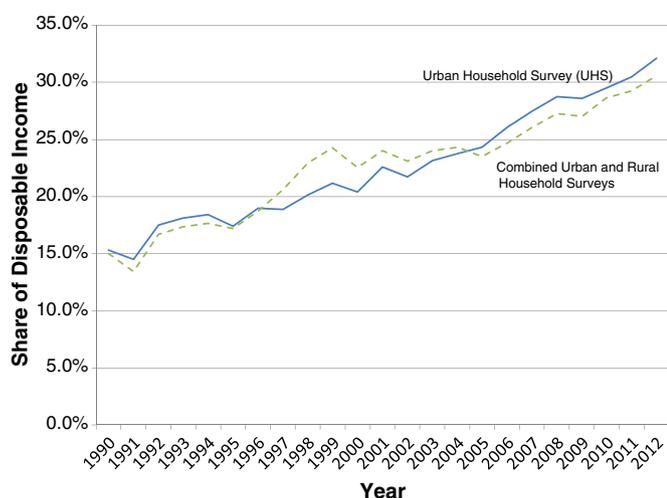


Fig. 1. Household saving rates.

Source: National Bureau of Statistics, Urban and Rural Household Surveys.

been a particularly pronounced increase in saving rates among households with younger and older household heads (see Fig. 2).³

Our main contribution to the literature on Chinese savings is to show that the rise in income uncertainty and the 1997 pension reform can together explain more than half of the observed rise in household savings rates as well as the dramatic shift in the age-savings profile. The existing literature analyzing the determinants of household savings in China has been largely focused only on the level or trend of the household savings rate. This literature includes papers using aggregate data (e.g., Kujis, 2006; Modigliani and Cao, 2004), provincial-level data (e.g., Horioka and Wan, 2007; Kraay, 2000; Qian, 1998) and micro data at the household or individual levels. Some of these studies find an important role for demographic considerations in explaining aggregate savings patterns. However, demographic variables tend to fare poorly when explaining household-level data (Chamon and Prasad, 2010). In a recent study mainly using provincial data, Wei and Zhang (2011) conclude that about half of the increase in household savings is related to imbalances in the sex ratio; households with male offspring save more in order to improve their marriage prospects. Banerjee et al. (2010) use a single-year cross-sectional survey to examine the effect of fertility on household savings. In work that is more closely related to ours, Song and Yang (2010) use household-level cross-sectional data and attribute much of the rise in household savings to expectations of a slowdown in income growth over the life cycle. One other study that looks at precautionary motives is that of Meng (2003), who uses cross-sectional household-level data to identify the effect of employment displacement on the consumption behavior of urban households and shows that savings help smooth those shocks.

Our initial contribution is to evaluate the effects of macroeconomic shifts on income uncertainty at the household level in China. We examine the evolution of household income using a sample of urban households tracked by the China Health and Nutrition Survey (CHNS) over the period 1989 to 2009. We exploit the panel aspect of the dataset to characterize the rise in income uncertainty and decompose the variance of income into components attributable to permanent versus temporary income shocks, following Meghir and Pistaferri (2004) and Blundell et al. (2008).⁴ We find strong trend growth in both the mean and the

³ This figure is based on a ten-province/municipality subsample of the full urban household survey data. We discuss this sample in more detail later in the paper. A similar pattern is also documented by Chamon and Prasad (2010) and Song and Yang (2010).

⁴ Also see Lillard and Weiss (1979), MaCurdy (1982), Abowd and Card (1989), Moffitt and Gottschalk (1995), and Baker and Solon (2003). Articles in a recent special issue of the *Review of Economic Dynamics* (2010) document the evolution of idiosyncratic income risk in different countries.

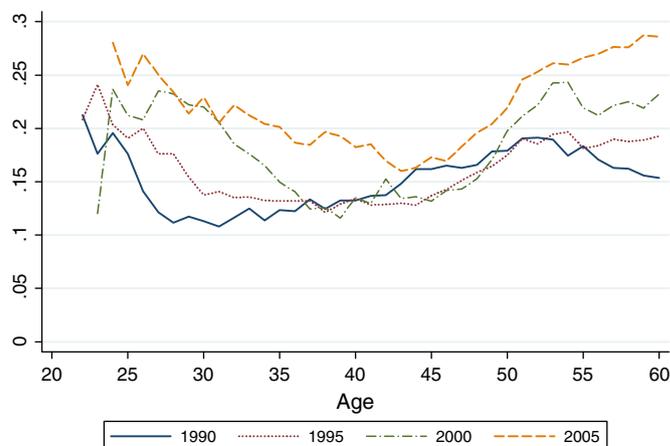


Fig. 2. Urban household saving rates by age of head. Notes: Based on a 10 province/municipality subsample of the National Bureau of Statistics Urban Household Survey. Saving rates smoothed by a moving average with 4 neighboring age averages. For details on the data, and how saving rates are defined, please refer to Chamon and Prasad (2010).

variance of total household income. We also document a substantial trend increase in the variance of transitory shocks to household income. There is also some evidence of an increase in the variance of permanent shocks, although this result is far less robust. This pattern is in line with a large literature on how technological and sectoral shifts and the associated labor reallocation can generate higher transitory uncertainty even if some of these shifts themselves are permanent.

Based on these results, we conduct a calibration of a simple buffer-stock/life-cycle model of savings to evaluate the implications of rising uncertainty on household saving rates, using the approach of Carroll (1997). We find that the rising variance of transitory shocks to income can help explain the rise in the savings of households with young household heads. For plausible parameter values, saving rates initially increase by over 4 percentage points for households with household heads in their twenties to mid-thirties. Since households with younger heads have a lower buffer stock of savings, an increase in transitory income variance causes them to save more in order to adjust their buffer stock to the riskier environment. After that initial adjustment, the response in saving rates declines over time (although young households entering the economy with no initial assets will continue to save 4 percentage points more than they would have had under the lower risk environment). Households with older household heads, which have already accumulated significant savings, can more easily accommodate transitory shocks.

To explain the increase in saving rates among households with older household heads, we turn to pension reform as a more promising explanation, calibrating the model to match changes in pension rules. Prior to the reform, urban workers received pensions through their employers—predominantly state-owned enterprises. These pensions had a replacement ratio of about 75–80% relative to the average wage (Dunaway and Arora, 2007; Sin, 2005). Workers retiring after 1997 are covered under the reformed system. They receive a social pension corresponding to 20% of the average local wage, the amount accumulated in individual retirement accounts and a supplementary “transition pension.” Sin (2005) estimates the replacement rate under different scenarios and concludes that, under the terms of the new pension rules, the replacement rate for the transition generation is around 60% of the average wage.⁵

⁵ The social pension is financed by employer contributions of 17% of wages. The individual accounts are financed by employer and employee contributions of 3% and 8% of wages, respectively (see Sin, 2005, and Dunaway and Arora, 2007, for more details). Herd et al. (2010) document that labor mobility is impeded by limited pension portability under the current system and also note that effective replacement rates are projected to decline further under the current rules.

Our calibration exercise suggests that a decline in the replacement rate from 75% to 60% of pre-retirement income can explain a 6–8 percentage point increase in saving rates for households whose heads are in their fifties and approaching retirement. As expected, the effect is more muted for households with younger household heads, who have a longer horizon to adjust their savings in response to the change in pension regime (the initial increase is one percentage point for households with heads in their thirties). But even the youngest cohorts end up saving 6 percentage points more by the time they are in their fifties than the pre-reform cohorts did.

In short, our calibration of a standard buffer-stock/life-cycle model of savings shows that higher income uncertainty and pension reforms can together explain much of the rise in average savings among urban households in China (as suggested by Blanchard and Giavazzi, 2006). Moreover, the calibrated response to saving rates implies changes to the cross-section of savings over time that are sharper among households at the two ends of the age distribution of household heads. Even 10 years after the initial increase in uncertainty and pension reform, we estimate the youngest and the oldest households save 5 percentage points more than before those changes, compared to only 3.5 percentage points more for those in their late thirties–early forties.

The results from our model better match the time profile of the changes in household savings when the reforms are modeled as a two-step process and households are assumed to learn about the shocks gradually. This alternative set-up implies a gradual increase in savings rates, which is more consistent with the evolution of urban household saving rates in China. The resulting increase in savings is back-loaded, with aggregate saving rates increasing by 5.6% in 10 years after the initial shock (about two-thirds of the increase observed in the aggregate urban data) and also exhibiting a U-shaped age-profile.

Our results are robust to alternative parameterizations of the model and conservative assumptions about the rise in transitory income uncertainty. Incorporating an increase in the variance of permanent shocks to income would strengthen the increase in aggregate savings, although we do not emphasize this result since our empirical analysis does not yield robust estimates of a shift in the variance of permanent shocks.

2. Dataset

We use data from the China Health and Nutrition Survey (CHNS).⁶ The survey is based on a multistage, random cluster process that yields a sample of about 4400 households with a total of 19,000 individuals that are tracked over time. The sample covers nine provinces that vary substantially in terms of geography, economic development, and other socioeconomic indicators.⁷ This survey was conducted in 1989, 1991, 1993, 1997, 2000, 2004, 2006 and 2009.

The sample in each province is drawn from a multistage random cluster process. Counties are stratified by income and a weighted sampling scheme is used to randomly select four counties in each province, in addition to the capital or main city, and a lower income city. The 1991 wave surveyed only individuals belonging to the original 1989 sample. In the 1993 wave, all new households formed from households in the previous survey sample were added to the sample. From the 1997 wave onwards, the sample includes newly formed households from

⁶ The survey is a collaborative effort between the Carolina Population Center at the University of North Carolina at Chapel Hill and the National Institute of Nutrition and Food Safety at the Chinese Center for Disease Control and Prevention. Details are at <http://www.cpc.unc.edu/projects/china>. Since it contains income data, the CHNS has been used to study income inequality (e.g., Li and Zhu, 2006) as well as other issues that require panel data, such as household income mobility (e.g., Ding and Wang, 2008).

⁷ The nine provinces are: Guangxi, Guizhou, Heilongjiang, Henan, Hubei, Hunan, Jiangsu, Liaoning, and Shandong.

the original sample, as well as additional households and new communities added to the sample to replace those households and communities that were no longer participating in the survey.

We use both individual and household files from the CHNS and focus on the urban sub-sample, consisting of households who do not have income from farming, raising livestock, fishing and gardening. The rural population exhibits much higher variance of earnings shocks (both permanent and transitory) relative to the urban population, probably due to the inherently more variable nature of agricultural incomes. Our baseline analysis involves a sample of households with household heads who are between the ages of 25 and 60, not a student, and for whom we have complete information on age and education. To avoid changes in household composition, we retain in the sample households whose heads remain the same. We include households in every year in which they appear in the data and satisfy these requirements. Our final sample is an unbalanced panel consisting of 1484 households.⁸ While this is a relatively small sample, the CHNS is the only publicly available panel dataset for Chinese households that can be used to quantify the variance and persistence of shocks to income. We focus on household income as it is more relevant than labor earnings when considering household consumption and savings decisions. Household income comprises labor income, private and public transfers, and subsidies, which constitute firm-level nonwage compensation to the worker and include subsidies on gas, food, education, and housing as well as allowances for children.

Table 1 shows the number of observations in each year and also presents summary statistics for the analysis sample. From 1989 to 2009, real mean annual household income more than triples, from 14,204 to 45,766 RMB at constant 2009 prices. Rising education levels in the population are reflected in the steadily increasing proportion of workers in our sample who have a high school degree (including a vocational training equivalent) or higher levels of education. The state-owned and collective enterprise (SOCE) sector—including government units, state-owned enterprises, and large collective enterprises (with a provincial or city government as the principal owner)—still plays an important role in the Chinese economy.⁹ In our sample, the proportion of workers employed in this sector falls from 82% in 1989 to 54% in 2009.¹⁰

3. A decomposition of permanent and transitory shocks to household income

In this section, we describe our methodology to decompose the variance in household income into the components attributable to permanent and temporary shocks. Following the literature modeling earning dynamics, we first run Mincerian income regressions that allow us to control year by year for cross-sectional income variation attributable to economy-wide shifts in the returns to observed household characteristics. In our preferred specification, we regress log income on region dummies (East, Northeast, Midwest and West), age and age squared, dummies for the education level of the household head (middle school or less, high school, some college), interactions

⁸ Attrition, introduction of new respondents into the survey, transitions into and out of employment, and aging affect households' and individuals' movement into and out of the analysis sample in different years.

⁹ State enterprise reform has involved selective privatization and hardening of budget constraints (reductions in explicit state subsidies) for the remaining enterprises. For more details on the reform process and how it has affected the operations and labor structure of these firms, see Lin et al. (1998), Bai et al. (2006) and Li and Putterman (2008). Brandt et al. (2008) analyze the effects of the reallocation of labor from the state sector to the non-state sector.

¹⁰ In the empirical work, we use characteristics of the household head. The head of the household as identified from the survey is typically a male (84% of the households in our sample), the primary earner (70% of households), and the oldest person in the household (68% of households).

Table 1
Summary statistics.

Wave	Observations (households)	Household Size Mean	Labor Earnings (in RMB at 2009 prices)		Income (in RMB at 2009 prices)		At least some high school		Head employed by SOCE	
			Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.
1989	473	3.8	8161.6	6045.6	14,204.4	8379.9	0.33	0.47	0.82	0.39
1991	446	3.6	8235.3	4374.0	13,100.2	6694.1	0.37	0.48	0.80	0.40
1993	352	3.5	10,003.3	7226.6	15,169.4	12,985.5	0.41	0.49	0.76	0.43
1997	416	3.3	13,803.4	10,241.7	17,836.5	11,647.1	0.52	0.50	0.69	0.46
2000	482	3.2	19,062.3	19,453.6	24,332.4	20,617.8	0.57	0.50	0.66	0.47
2004	494	2.9	24,298.4	19,403.1	32,786.8	27,418.6	0.64	0.48	0.59	0.49
2006	476	2.9	29,205.5	32,626.5	36,924.1	37,689.4	0.68	0.47	0.55	0.50
2009	541	2.9	36,543.6	39,237.7	45,766.1	55,797.0	0.59	0.49	0.54	0.50

Notes: Based on an unbalanced panel of urban households from the China Health and Nutrition Survey, with household heads who are aged 25–60, not a student, with complete age and education information, and not reporting positive income from farming and raising livestock.

between age and education dummies, and marital status. This regression is run separately for each year.¹¹

Our focus in this paper is on household-specific income uncertainty, so we will mostly work with residuals from the first-stage regressions. In effect, we analyze within-group variations in income that cannot be explained by the household characteristics included in those regressions. We use the residuals to estimate the permanent and transitory components of income:

$$\begin{aligned} y_{iat} &= u_{iat} + v_{iat} \\ u_{iat} &= u_{i,a-t-1} + \omega_{iat} \end{aligned} \tag{1}$$

where y_{iat} is the log earning residuals for household i with household head aged a in year t from the Mincerian regression, u_{iat} is the permanent component, and v_{iat} is the transitory component including measurement error. The assumption that the permanent component is a random walk is standard in the literature and is consistent with our data as we find negligible autocovariances of income growth beyond the first lag. Since the income regressions are run separately for each year, the residuals correspond to within-group measures of log income, taking out the mean effects of region, education level, age and the other household characteristics that we have controlled for. The permanent shocks ω and transitory shocks v to earnings have zero means and are mutually orthogonal. They are i.i.d. across household, time and age.¹² We assume:

$$\begin{aligned} \text{var}(\omega_{iat}) &= \sigma_{\omega t}^2 \\ \text{var}(v_{iat}) &= \sigma_{\xi t}^2 \end{aligned} \tag{2}$$

In other words, the variances of permanent and transitory shocks change by year but do not depend on age. This in effect amounts to

¹¹ To conserve space, we do not report the regression results in detail here. The estimates show rising returns to education. The pattern of returns to potential labor market experience is less clear. We re-estimated the income regressions using alternative sets of covariates and also tried using the detrended log of total household income. The trends in estimated transitory and permanent income uncertainty that we report below remain very similar. See Table 2 discussed in the next section for details.

¹² The transitory shocks do not appear to be serially correlated. We estimate the following autocovariances of unexplained income growth at lags 1 to 3 (standard errors in parentheses): -0.142 (0.016), 0.001 (0.017), and 0.002 (0.018). Autocovariances of order 2 and higher are not statistically significant. If we test the null hypothesis of zero autocovariances in income growth (allowing autocovariances to be different across years), we reject the null hypothesis at lag one but not for higher order lags. These results indicate that it is reasonable to assume that the permanent component is a random walk and that the transitory shocks are either i.i.d. or follow an MA(1) process. This is consistent with much of the literature (Abowd and Card, 1989; Blundell et al., 2008; Meghir and Pistaferri, 2004). Because of the gaps between years of observations in the data, it is not possible to further test the stochastic process of transitory shocks. As we discuss later, the permanent uncertainty identified by our model is consistent regardless of whether the transitory shock follows an MA(1) process or is i.i.d.

averaging over households with different ages (or in different cohorts) in each year.¹³ Later, we will examine how these variances differ across age groups. From here on, the subscript a will be dropped. The parameters to be estimated are: $\sigma_{\omega t}^2$ and $\sigma_{\xi t}^2$ for each survey wave: $t = \{1989, 1991, 1993, 1997, 2000, 2004, 2006, 2009\}$.

Suppose we observe household income in consecutive years. We can write one-year changes in income, Δy_{it} , as:

$$\begin{aligned} \Delta y_{it} &= y_{it} - y_{it-1} = \omega_{it} + v_{it} - v_{it-1} \\ \Delta y_{it-1} &= y_{it-1} - y_{it-2} = \omega_{it-1} + v_{it-1} - v_{it-2} \end{aligned} \tag{3}$$

where y_{it} is the log earning residual for household i in year t from the Mincerian regression, ω_{it} is the permanent shock, and v_{it} is the transitory shock including measurement error. Under the assumptions that the permanent shocks and transitory shocks are orthogonal to each other and i.i.d. over time, the variance and autocovariance of one-year changes in income are as follows:

$$\begin{aligned} \text{cov}(\Delta y_{it}, \Delta y_{it-1}) &= \text{cov}(\omega_{it} + v_{it} - v_{it-1}, \omega_{it-1} + v_{it-1} - v_{it-2}) \\ &= -\text{var}(v_{it-1}) = -\sigma_{\xi t-1}^2 \\ \text{var}(\Delta y_{it}) &= \text{var}(\omega_{it} + v_{it} - v_{it-1}) = \sigma_{\omega t}^2 + \sigma_{\xi t}^2 + \sigma_{\xi t-1}^2 \end{aligned} \tag{4}$$

This implies that the transitory variance in year t can be estimated from the one-period lagged autocovariance of income changes, i.e., the covariance between income changes at time t and income changes at time $t - 1$. The variance of permanent shocks in year t can be identified from:

$$\begin{aligned} \sigma_{\omega t}^2 &= \text{var}(\Delta y_{it}) - \sigma_{\xi t}^2 - \sigma_{\xi t-1}^2 = \text{var}(\Delta y_{it}) + \text{cov}(\Delta y_{it}, \Delta y_{it-1}) \\ &\quad + \text{cov}(\Delta y_{it+1}, \Delta y_{it}) \end{aligned} \tag{5}$$

In other words, the permanent variance is the variance of income changes in year t , stripping off the contribution from the transitory variances in t and $t - 1$.

With four years of data $\{t + 1, t, t - 1, t - 2\}$, we would be able to identify $\sigma_{\omega t}^2$, $\sigma_{\omega t-1}^2$, $\sigma_{\xi t}^2$. Note that the parameters are identified nonparametrically without making any distributional assumptions about the shocks. Nor does the identification involve any assumption about $\sigma_{\omega 0}^2$, the initial variance of permanent earnings. This is an important advantage over alternative identification strategies (e.g., moments using earning levels), particularly for a fast-growing economy where $\sigma_{\omega 0}^2$ is likely to be nonstationary.

The uneven spacing of the CHNS waves complicates the analysis since we need to use n -year rather one-year income changes. Applying

¹³ We focus on the year effect and therefore the age and cohort effects cannot be separated. Given our sample size, we cannot allow variances to also vary by age (or cohort).

the same formulas as above, we derive the permanent and transitory variances from the CHNS data as follows:

$$\begin{aligned}
 \text{cov}(Y_{i,93} - Y_{i,91}, Y_{i,91} - Y_{i,89}) &= -\sigma_{\xi 91}^2 \\
 \text{cov}(Y_{i,97} - Y_{i,93}, Y_{i,93} - Y_{i,91}) &= -\sigma_{\xi 93}^2 \\
 \text{cov}(Y_{i,00} - Y_{i,97}, Y_{i,97} - Y_{i,93}) &= -\sigma_{\xi 97}^2 \\
 \text{cov}(Y_{i,04} - Y_{i,00}, Y_{i,00} - Y_{i,97}) &= -\sigma_{\xi 00}^2 \\
 \text{cov}(Y_{i,06} - Y_{i,04}, Y_{i,04} - Y_{i,00}) &= -\sigma_{\xi 04}^2 \\
 \text{cov}(Y_{i,09} - Y_{i,06}, Y_{i,06} - Y_{i,04}) &= -\sigma_{\xi 06}^2 \\
 \text{var}(Y_{i,93} - Y_{i,91}) &= \sigma_{\omega 93}^2 + \sigma_{\omega 92}^2 + \sigma_{\xi 93}^2 + \sigma_{\xi 91}^2 \\
 &= \sigma_{\omega 93}^2 + \sigma_{\omega 92}^2 - \text{cov}(Y_{i,97} - Y_{i,93}, Y_{i,93} - Y_{i,91}) \\
 &\quad - \text{cov}(Y_{i,93} - Y_{i,91}, Y_{i,91} - Y_{i,89}) \\
 &\quad \dots \\
 \text{var}(Y_{i,04} - Y_{i,00}) &= \sigma_{\omega 04}^2 + \sigma_{\omega 03}^2 + \sigma_{\omega 02}^2 + \sigma_{\omega 01}^2 + \sigma_{\xi 04}^2 + \sigma_{\xi 00}^2 \\
 \text{var}(Y_{i,06} - Y_{i,04}) &= \sigma_{\omega 06}^2 + \sigma_{\omega 05}^2 + \sigma_{\xi 06}^2 + \sigma_{\xi 04}^2.
 \end{aligned} \tag{6}$$

Each of the equations is a moment condition—the left hand side is a variance or covariance estimated directly from the data and the right hand side is a function of parameters implied by the income process. The permanent variance can still be seen as identified by the variance of income changes between year $t + k$ and year t , stripping off the contribution from the transitory variances in years $t + k$ and t . We are able to identify six years of the transitory income risk, all except 2009 and 1989. We do not make any assumption about the transitory variances in those two years and, hence, we are able to identify five permanent variances:

$$\begin{aligned}
 &\sigma_{\omega 93}^2 + \sigma_{\omega 92}^2 \\
 &\sigma_{\omega 97}^2 + \sigma_{\omega 96}^2 + \sigma_{\omega 95}^2 + \sigma_{\omega 94}^2 \\
 &\sigma_{\omega 00}^2 + \sigma_{\omega 99}^2 + \sigma_{\omega 98}^2 \\
 &\sigma_{\omega 04}^2 + \sigma_{\omega 03}^2 + \sigma_{\omega 02}^2 + \sigma_{\omega 01}^2 \\
 &\sigma_{\omega 06}^2 + \sigma_{\omega 05}^2.
 \end{aligned} \tag{7}$$

In practice, the permanent and transitory variances are estimated jointly in one step. We estimate the model using an equally-weighted minimum distance estimator, a standard approach in the literature since Moffitt and Gottschalk (1995). The model is just identified.

Note that our estimated variance of transitory shocks could be biased upwards for two reasons. One is that in light of classical measurement errors (i.i.d.), the estimated variance of transitory shocks will be inconsistent and biased upwards. This should not drive the trend in transitory variance unless the variance of measurement errors itself has a trend. Second, we cannot exclude the possibility that the transitory shocks may follow an MA(1) process (see footnote 9), implying that the identified transitory variance also includes the transitory shocks from the previous year. However, without additional assumptions, it is not possible to identify the MA(1) process given the gaps between our sampling years. Since we are looking at n -year differences ($n \geq 2$), even if $v_{it} = \xi_{it} + \theta \xi_{it-1}$ (workers take two years to recover from a transitory shock to income), our estimates of the variance of permanent shocks are still consistent. To see this:

$$\begin{aligned}
 \text{cov}(Y_{i,93} - Y_{i,91}, Y_{i,91} - Y_{i,89}) &= -\sigma_{\xi 91}^2 - \theta^2 \sigma_{\xi 90}^2 \\
 \text{cov}(Y_{i,97} - Y_{i,93}, Y_{i,93} - Y_{i,91}) &= -\sigma_{\xi 93}^2 - \theta^2 \sigma_{\xi 92}^2 \\
 \text{var}(Y_{i,93} - Y_{i,91}) &= \sigma_{\omega 93}^2 + \sigma_{\omega 92}^2 + (\sigma_{\xi 91}^2 + \theta^2 \sigma_{\xi 90}^2) + (\sigma_{\xi 93}^2 + \theta^2 \sigma_{\xi 92}^2) \\
 &= \sigma_{\omega 93}^2 + \sigma_{\omega 92}^2 - \text{cov}(Y_{i,93} - Y_{i,91}, Y_{i,91} - Y_{i,89}) \\
 &\quad - \text{cov}(Y_{i,97} - Y_{i,93}, Y_{i,93} - Y_{i,91}).
 \end{aligned} \tag{8}$$

In order to account for these two potential biases, when calibrating the savings model we will assume that the true transitory uncertainty is only one-half of the transitory variance actually identified from our

estimates. This is a rather conservative assumption. Researchers using U.S. household income data typically find that the estimated MA(1) parameter for the transitory shock is small (between -0.1 and -0.2 ; see, e.g., Blundell et al., 2008). So the upward bias of the estimated transitory uncertainty due to serial correlation of the transitory shocks ($\theta^2 \sigma_{\xi t-1}^2$) is likely to be small.

That leaves the potential bias from measurement errors. If we assume that measurement error accounts for half of the identified transitory variance, then our estimates imply that measurement error alone would explain more than 40% of the variance of income growth. In fact, researchers conducting validation studies using U.S. data find that measurement error accounts for only around a quarter of the variance of growth of earnings.¹⁴ It is worth emphasizing that, so long as the variance of the measurement error itself does not have a trend, our estimates of the trends in the variances of transitory and permanent shocks are still consistent.

4. Earnings decomposition results

Table 2 reports, in panels A and B respectively, estimates of the variances of the permanent and transitory shocks to household income and earnings over time. Standard errors are computed using a block bootstrap procedure. One should bear in mind that the sample size is small, which limits the power of statistical inference we are able to obtain (using the only available panel dataset for the question we are interested in). The first column of panel A shows that, for the full urban sample, there is no clear trend in the variance of permanent shocks to income. While the point estimate increases from 0.012 to 0.030 from 1993 to 1997, the difference is not statistically significant at the 5% level. The point estimate increases further to 0.043 in 2006, but the difference relative to the estimate for 1993 is also not statistically significant.

Columns 2–7 provide estimates based on alternative specifications for the Mincerian regression. In column 2, we drop the age * education interaction terms. In column 3, we replace household size fixed effects with fixed effects over the number of income earners. In column 4, we use only a constant as a regressor. In column 5, we trim the sample to exclude households in the top and bottom 1% of the distribution of raw household incomes. Finally, in columns 6 and 7, we restrict the sample to households headed by workers aged 30–60 years and 25–55 years, respectively. The results remain broadly the same across these specifications, namely that there is no clear and statistically significant trend in the variance of permanent shocks to income.

In panel B of Table 2, we present estimates of the variance of transitory shocks to household income. The point estimates in column 1 rise from 0.061 in 1991 to 0.181 by 2006, and this increase is statistically significant at the 1% level.¹⁵ In other words, income uncertainty due to transitory shocks to income almost triples from the beginning of the 1990s to the 2000s. A similar pattern holds across the different specifications estimated in columns 2–7.

Table 3 presents the estimates from our baseline specification across different sub-samples. As in the previous table, panel A reports the estimates for the permanent variance. Column 1 reports the estimates for the whole sample (and is identical to column 1 of Table 2). Column 2 reports the estimates for a sample of households whose head worked in the SOCE sector when the household entered the panel.¹⁶ The results are similar to those of column 1, and again do not suggest any trend (and differences in point estimates between these two columns are

¹⁴ See Bound and Krueger (1994). In the case of non-classical measurement errors, Pischke (1995) finds that the transitory variance is less contaminated due to the negative correlation of measurement errors with transitory earnings.

¹⁵ The jump in transitory variance from 1991 to 1993 largely reflects a higher variance of raw log income in 1993 that partially settles back down in 1997 (this can also be seen in the jump in the standard deviation of household income in 1993 in Table 1).

¹⁶ The results are similar if we define the SOCE subsample based on SOCE employment throughout the survey (i.e., excluding workers who start in the SOCE sector but later move to the non-SOCE sector).

Table 2
Estimated variances of permanent and transitory shocks to urban household income.

Year	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>A. Permanent variance</i>							
1993	0.012 (0.012)	0.012 (0.012)	0.001 (0.009)	0.026 (0.007)	0.011 (0.010)	0.022 (0.012)	0.024 (0.016)
1997	0.030 (0.011)	0.031 (0.011)	0.028 (0.012)	0.025 (0.031)	0.016 (0.010)	0.029 (0.011)	0.023 (0.012)
2000	0.030 (0.018)	0.032 (0.018)	0.027 (0.016)	0.040 (0.020)	0.038 (0.014)	0.019 (0.016)	0.040 (0.018)
2004	0.031 (0.014)	0.031 (0.014)	0.021 (0.011)	0.039 (0.018)	0.022 (0.008)	0.028 (0.014)	0.020 (0.013)
2006	0.043 (0.025)	0.041 (0.025)	0.036 (0.021)	0.055 (0.045)	0.028 (0.020)	0.048 (0.025)	0.058 (0.036)
<i>B. Transitory variance</i>							
1991	0.061 (0.016)	0.061 (0.016)	0.083 (0.019)	0.068 (0.020)	0.054 (0.014)	0.053 (0.016)	0.051 (0.017)
1993	0.154 (0.030)	0.154 (0.030)	0.163 (0.030)	0.162 (0.039)	0.124 (0.024)	0.136 (0.032)	0.141 (0.034)
1997	0.101 (0.047)	0.101 (0.047)	0.072 (0.041)	0.091 (0.053)	0.114 (0.037)	0.110 (0.040)	0.107 (0.048)
2000	0.174 (0.053)	0.172 (0.052)	0.163 (0.045)	0.163 (0.038)	0.082 (0.025)	0.184 (0.061)	0.148 (0.045)
2004	0.192 (0.035)	0.197 (0.035)	0.154 (0.025)	0.200 (0.051)	0.167 (0.025)	0.178 (0.031)	0.218 (0.036)
2006	0.181 (0.045)	0.184 (0.046)	0.160 (0.040)	0.168 (0.090)	0.176 (0.037)	0.173 (0.048)	0.168 (0.065)

Notes: Variance estimates based on the decomposition described in Section 3. Standard errors are computed using a block bootstrap procedure, and are reported in parenthesis. Column (1) corresponds to our preferred specification, where residuals are based on a Mincerian regression with age, age squared, education dummies, interactions of age and education dummies, region dummies, household size dummies, and marital status. Columns 2–7 consider alternative specifications for that Mincerian regression. Column (2) drops the interactions of age and education dummies. Column (3) replaces household size dummies with dummies for the number of income earners. Column (4) is based on a specification with no controls (only the constant as a regressor). Column (5) trims out the top and bottom 1% of the sample after sorting by raw household income. Columns (6) and (7) are based on a sample with workers aged 30–60 and 25–55 years old respectively.

not statistically significant). When we split the sample by birth cohort (head born before or after 1955), the results remain broadly similar, suggesting no clear trend. Splitting the sample by educational attainment of the household head (with or without high school degree) yields rather noisy results, with large standard errors for the group of households with less-educated household heads. This is in part driven by the large increase in education levels over the sample (households with less educated heads are concentrated in the initial waves and those with more educated heads in the most recent waves). But again, the estimates do not suggest a trend, although they do suggest that the permanent uncertainty facing households with less-educated heads is larger than the permanent uncertainty facing households with more educated heads.

Turning to the transitory variance (panel B), the results are similar in the full and SOCE samples (columns 1 and 2), although the increase seems more gradual for the latter. The subsamples where we divide observations by cohort and educational attainment of the household head have noisier patterns, but are generally consistent with the trend of a substantial increase in the variance of transitory shocks since the early 1990s. The pattern of a trend increase in transitory uncertainty remains if we exclude transfers and subsidies from household income (last column).¹⁷ In this case, the estimated level of uncertainty is generally higher in most years compared to the level for total household income, consistent with the prior that transfers and subsidies serve as partial insurance against idiosyncratic shocks to household income.

¹⁷ In the early stages of reform, SOCEs offered workers higher levels of subsidies to compensate for noncompetitive wages and then reduced them as their budget constraints were tightened (reduced transfers from the state to SOCE firms). The ratio of subsidies to total compensation was as high as 35% in the 1991 wave, but steadily declines to about 5% in the 2006 wave.

Table 3
Estimated variances of permanent and transitory shocks to urban household income across different samples.

Year	Family income						Family income excl. trans. and sub.
	All	SOCE	Younger cohort	Older cohort	Less edu.	More edu.	All
<i>A. Permanent variance</i>							
1993	0.012 (0.012)	0.019 (0.010)	0.033 (0.045)	0.006 (0.010)	0.015 (0.018)	0.000 (0.008)	0.015 (0.015)
1997	0.030 (0.011)	0.021 (0.008)	0.054 (0.059)	0.037 (0.016)	0.055 (0.032)	0.008 (0.010)	0.039 (0.016)
2000	0.030 (0.018)	0.013 (0.014)	0.017 (0.024)	0.022 (0.027)	0.072 (0.045)	0.008 (0.014)	0.035 (0.017)
2004	0.031 (0.014)	0.013 (0.009)	0.016 (0.012)	0.066 (0.037)	0.096 (0.055)	0.012 (0.010)	0.040 (0.014)
2006	0.043 (0.025)	0.019 (0.021)	0.057 (0.032)	0.022 (0.035)	0.075 (0.085)	0.039 (0.028)	0.044 (0.023)
<i>B. Transitory variance</i>							
1991	0.061 (0.016)	0.034 (0.010)	0.059 (0.029)	0.066 (0.018)	0.066 (0.027)	0.062 (0.021)	0.080 (0.021)
1993	0.154 (0.030)	0.092 (0.021)	0.202 (0.109)	0.124 (0.027)	0.173 (0.062)	0.135 (0.032)	0.185 (0.035)
1997	0.101 (0.047)	0.089 (0.027)	0.175 (0.097)	0.067 (0.044)	0.082 (0.079)	0.112 (0.039)	0.115 (0.042)
2000	0.174 (0.053)	0.183 (0.055)	0.152 (0.055)	0.199 (0.105)	0.125 (0.101)	0.205 (0.062)	0.168 (0.047)
2004	0.192 (0.035)	0.192 (0.032)	0.212 (0.043)	0.140 (0.063)	0.098 (0.077)	0.210 (0.043)	0.188 (0.037)
2006	0.181 (0.045)	0.176 (0.054)	0.187 (0.058)	0.180 (0.085)	0.249 (0.123)	0.169 (0.057)	0.185 (0.045)

Notes: Variance estimates are based on the decomposition described in Section 3. “All” refers to the full sample. SOCE subsample restricted to those households whose head was an SOCE employee when the household first appeared in the sample. Younger cohort and older cohort households are defined as those whose household heads were born after and before 1955, respectively. “Less educated” is the group of households whose head does not have a high school degree. The last column considers the full sample but using an income measure that excludes subsidies and transfers (both public and private).

What accounts for the rising variance of transitory income shocks experienced by Chinese households? While building a structural model to explain these facts is beyond the scope of this paper, we provide some descriptive evidence from labor market turnover. A number of papers have documented that higher labor market turnover (both job to job transitions and transition into and out of unemployment) could lead to higher transitory uncertainty (see, e.g., Topel and Ward, 1992; Gottschalk and Moffitt, 1994). Gottschalk and Moffitt (2009) find that the rise in transitory variance explains about half of the rise in cross-sectional income inequality in the U.S. through the late 1980s and that this increase in earnings instability is in part attributable to greater instability in jobs and higher labor market turnover.

Fig. 3 shows that in urban China the transition rate from employment to unemployment for all workers increases sharply in the late 1990s and continues to rise during the 2000s, corresponding to the period when our estimates suggest that transitory income uncertainty began rising.¹⁸ The transition from employment in the SOCE sector to employment in the non-SOCE sector is also high starting in the mid-1990s, whereas the transition from non-SOCE employment to SOCE sector employment has fallen. In addition to these labor market outcomes, the transition from a centrally planned economy to a market economy may have resulted in an increase in firm-level volatility related in part to state enterprise restructuring and a tighter link between wages and firm-level performance. Wages paid to workers may be increasingly tied to firm performance and more reflective of individual productivity due to tightening of budget constraints on SOCEs, increased competition

¹⁸ We define an individual as unemployed if he/she does not receive any wage or business income in a given year.

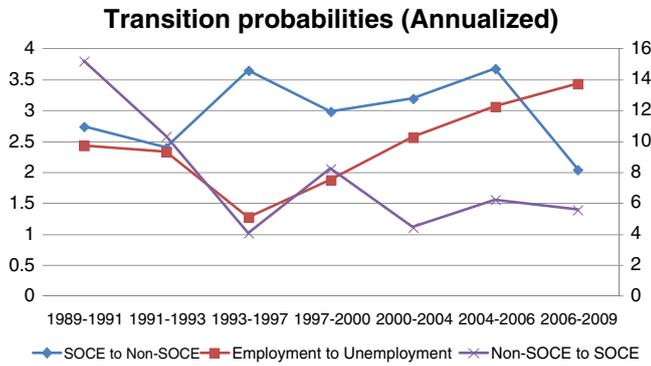


Fig. 3. Labor market turnover: annualized transition probabilities. Notes: Transition rates annualized by taking the Markov transition matrix between two surveys to the power of $1/n$, where n is the number of years between the survey pair.

and more openness to foreign trade (see Benson and Zhu, 1999; Gang et al., 1998; Groves et al., 1994).

Comin et al. (2009) show that firm-level instability increased after 1980 in the U.S. (particularly for large firms with volatile sales), corresponding to a period of higher transitory variance of labor earnings documented in the U.S. Violante (2002) shows that skill-neutral technological change could result in an increase in the variance of the transitory component of earnings. In his model, workers learn vintage-specific skills and, when separating from their jobs, can only partially transfer their skills across machines. Therefore, technological acceleration reduces skill transferability and increases wage losses upon separation, which can increase cross-sectional wage variability in an economy undergoing major technological shifts and/or significant labor market churning. The rate of technological change in China since the 1990s has been even faster than in the U.S., due to the transition process and catching-up effects. This makes skill-biased technological change a promising candidate to help explain the increase in the variance of transitory income shocks.

5. Implications of the shifts in income variance for precautionary savings

Greater uncertainty in earnings at the microeconomic level can have macroeconomic implications. One important channel is the impact of greater household-specific uncertainty on precautionary savings. In the absence of a strong social safety net and an underdeveloped financial system, this could lead households to self-insure by increasing their savings (Blanchard and Giavazzi, 2006; Chamon and Prasad, 2010). In order to quantify the effects of this rise in uncertainty on individual and aggregate savings, we now undertake a calibration of a precautionary savings model, building on the work of Carroll (1997) and Gourinchas and Parker (2002).¹⁹ This enables us to provide a quantitative measure of how the increase in the variance of transitory shocks to household income can translate into the rise in savings among the younger households observed in the data, while changes in pension rules can help explain the savings of the older households.

5.1. Stylized facts

To motivate this exercise, we turn again to Figs. 1 and 2. We will focus on 1997 as the base year for comparing saving rates, as it is the

¹⁹ See also Fuchs-Schündeln (2008) and Kaplan and Violante (2010). Our calibration exercise sets only a lower bound on the degree of precautionary saving attributable to earnings uncertainty. We consider the variance of different shocks to earnings only for workers who report positive earnings in each period. For workers who in reality face unemployment and the prospect of zero earnings, the precautionary savings motive could be even stronger.

year the pension reforms were instituted, and is also around the time when the variance in the estimated income process begins to rise. The average saving rate in the UHS increased by only 0.8 percentage points from 1993 to 1997. But it increased by 2.8 percentage points from 1997 to 2002, and by an additional 5.8 percentage points from 2002 to 2007. Ideally, we would like to evaluate the model's predictions relative to the observed increase in savings for a comparable sample of the UHS that covers only households with heads aged 25–60. This is the age group for which the income process was estimated in the previous section and for which the model is calibrated in the section below.

However, we do not have access to the full disaggregated UHS (or to the tabulations that would allow such calculation), and no other survey provides income and consumption expenditure data spanning the period studied in this paper. In the subsample of ten provinces/municipalities used in Chamon and Prasad (2010), the saving rates among 25–60 year olds increased by 1.1 percentage points from 1993 to 1997, which is reasonably close to the increase for the UHS as a whole during that period.²⁰ From 1997 to 2005 (the last year covered in that subsample), the saving rate among 25–60 year olds increases by 4.8 percentage points, while the increase for the UHS as a whole was 5.4 percentage points. Given the close similarity in the trends and the fact that the full UHS is available for more recent years, we chose to use the full UHS as a basis for comparing the model predictions with data. This longer time coverage is particularly important when we consider gradual learning (Section 6.2).

In addition to the aggregate increase in the saving rate, we also aim to explain the evolution of its age-profile. For that we turn back to Fig. 2, which plots household saving rates as a function of the age of the head of household observed in the actual data for different years, based on the subsample of 10 provinces/municipalities used in Chamon and Prasad (2010). In the early 1990s, the age-savings profile in China was fairly typical of those in other economies, with saving rates increasing with age. Over time, saving rates have increased across the board. But the increase is particularly pronounced for households with relatively young household heads (those in their twenties and early thirties) and older household heads (aged in the mid-fifties and up). Consequently, by 2005 the age-savings profile has an unusual U-shaped pattern. Therefore, any empirically relevant explanation for the increase in savings rates must account not only for the substantial average increase, but also for the unusual way in which that increase was concentrated among the younger and older households. Our calibration below is able to capture these empirically relevant features.

5.2. Model and calibration

We assume an instantaneous CRRA utility function, with individuals maximizing the expected discounted flow of utility subject to a no-borrowing constraint:

$$\max \sum_t \beta^t \left(\prod_{j=0}^t s_j \right) E_t \left[\frac{C_t^{1-\gamma}}{1-\gamma} \right] \quad (9)$$

$$\text{s.t. } A_{t+1} = (1+r)(A_t + Y_t - C_t), \quad A_t \geq 0, \quad \forall t \quad (10)$$

where β is the discount factor, s is an age-dependent survival probability, C_t is the level of consumption in period t , γ is the coefficient of relative risk aversion, A_t is the level of assets, and Y_t represents income at time t . We assume that income is based on the same process estimated in

²⁰ The sample covers the following provinces: Anhui, Beijing, Chongqin, Ganshu, Guangdong, Hubei, Jiangsu, Liaoning, Shanxi, and Sichuan. Three of these overlap with the CHNS sample (Hubei, Jiangsu, and Liaoning).

the previous section for the working years, but permanent income becomes deterministic in the retirement period R at a particular fraction of the pre-retirement permanent income. That is:

$$\begin{aligned} y_t &= u_t + v_t & \text{if } t \leq R \\ u_t &= u_{t-1} + \omega_t \end{aligned} \quad (11)$$

$$\begin{aligned} y_t &= \eta u_t & \text{if } t > R \\ u_t &= u_{t-1} \end{aligned} \quad (12)$$

where y_t is $\log(Y_t)$, u_t is permanent income, v_t is transitory income, and ω_t is the shock to permanent income. The model is solved backwards starting from the last period of life using the endogenous grid point method developed by Carroll (2006). We calibrate the model assuming that working life begins at age 25, with an initial level of wealth of zero and initial level of permanent income equal to one. The discount factor β is 0.97. The real interest rate is 1.4% per annum, which matches the average real interest rate in China over the period 1989–2006.²¹ The coefficient of relative risk aversion γ is 4.5. We assume that economic agents live with certainty until the retirement age of 60, have a survival probability until age 85, and die with certainty if still alive at age 85.²² There are no bequests (for an individual who dies prior to age 85 with a positive level of assets, those assets are “lost”).²³

Permanent income in the retirement period is initially set such that $\eta = 75\%$ of pre-retirement permanent income, in line with the replacement rate prior to the 1997 reform. When we model the effects of the pension reform (which affects workers retiring after 1997), we will set $\eta = 60\%$ of pre-retirement permanent income.²⁴ In practice, there was significant heterogeneity in the manner in which the pension reforms were applied to different groups of workers. For instance, transitional arrangements partially protected some older workers from the full decline in the replacement ratio. Moreover, while the central government sets the policy framework for the new pension scheme, its administration was left to the discretion of local authorities, with pooling being done at the provincial or county/municipal level. This often resulted in large discrepancies on contribution rates and pensions between workers in different localities. Modeling this heterogeneity is beyond the scope of this paper, although uncertainty about the actual replacement rate can be viewed in our framework as another element of income uncertainty.²⁵

To calibrate the income process, we use the deterministic life-cycle growth rate of earnings in the CHNS sample. We regress the log of family income on a complete set of cohort dummies, household size, and a fourth-order polynomial in the age of the household head. We calculate the marginal effect of age on household income at each age.²⁶ The

predicted annual income growth is about 7% for the young, then ranges between 6 and 7% throughout most of the remaining work life before gradually declining to 2% close to retirement age.²⁷

We want to model how saving rates respond to changes in income uncertainty along the lines suggested by our empirical estimates in the previous section. We focus on family income and set the variance of permanent shocks to income at a constant level of 0.02, while the variance of transitory shocks goes from 0.03 in the baseline case up to 0.09. These variances are lower than the point estimates reported in the earlier section on account of the conservative assumption we make that half of the variance estimated in our empirical work is due to measurement error. This assumption reduces the effect of rising uncertainty on saving in our calibrations. Moreover, by not considering changes to the variance of permanent shocks, given the lack of a clear trend in the empirical estimates, our calibration results provide a conservative assessment of how rising uncertainty has affected savings.

6. Calibration results

We now present results from the calibration of the model described in the previous section. We consider two main experiments. The first is a one-time change in transitory variance in pension reforms. This illustrates the main mechanisms in the model. We then develop a more realistic variant with the pension reforms and the increase in transitory uncertainty split into a two-step change. We also allow for gradual learning about these shocks. With these features, the model does a better job of matching up to the temporal evolution of household saving rates.

6.1. One-off change

Fig. 4A plots the simulated age profile of the saving rate. We construct the age-savings profile by simulating the model for 10,000 households, and averaging their saving rates at each age. The dashed line corresponds to the profile of savings under the initial baseline assumptions about the variance of income. Consistent with this type of buffer-stock/life-cycle model, saving rates show a U-shaped pattern when plotted against age. Saving rates initially decline with age, since households with the youngest household heads typically start their working life cycle with no assets and need to save more in order to quickly build an adequate buffer stock of savings. Once that buffer stock is built, savings remain relatively low until the late thirties/early forties when earnings increase and life-cycle motives lead to a sharp increase in the saving rate.

The additional lines in this figure correspond to the age-savings profile after the change in the income process. Each line corresponds to the saving behavior that would result if the regime switch would occur starting at a given age of the household head (e.g. 25, 30, ..., 55), and after the initial jump we trace the behavior that would occur through the rest of the life cycle under those parameters. That change is more easily seen in Fig. 4B, which plots the change in the savings rate after the shock as a function of the age of the household head. If a household head were to begin working life at age 25 already under the higher uncertainty regime, that household would save about 4.5 percentage points more to begin with. The difference in saving rates relative to the baseline regime declines with age. For example, the initial jump for a household with a forty year old head is only about 2.5 percentage points. The reason for this pattern is the lower buffer stock of savings of the youngest households (since they start life with no buffer stock of savings). A lower buffer stock causes households with younger heads to respond more strongly to the shock to the transitory variance of income. The effect on households with older heads is

²¹ The real interest rate is based on the nominal interest rate on one-year bank deposits deflated by the annual CPI inflation rate.

²² Corecidence of different generations in one household could have implications for our analysis of life cycle savings. Saving for retirement is less important if one expects to live with and be supported by one's children during old age. By the same token, workers may have lower savings during their working years because of the burden of caring for elderly parents. Based on the trends in living arrangements, it is reasonable to assume that pensions and savings will provide the main source of support during old age, especially for younger cohorts, given that coresidence is likely to be less common in the future.

²³ For simplicity we assume a Poisson death process, calibrated to match life expectancy in China in 2009 (73.5 years). This results in a constant survival probability of 0.925 between t and $t + 1$ after retirement.

²⁴ The replacement rate should decline over time, given the nature of the pension formula. Sin (2005) projects the replacement rate for a male retiring at age 60 to decline to about 60, 55 and 50% of the average wage by 2010, 2020 and 2030, respectively. Thus, our assumption for the decline in the replacement rate is a conservative one, particularly for younger workers.

²⁵ Feng et al. (2011) discuss the pension reforms in detail. Their empirical analysis leads them to conclude that, consistent with our results, the reforms increased household saving rates.

²⁶ This assumes that there is no cohort effect on the growth rate of earnings. That is, while younger cohorts are much richer than older ones, the age profile of income growth is the same for both. One could argue that younger cohorts should expect slower growth as China's growth rate may eventually moderate.

²⁷ More precisely, the predicted income growth is on average: 6.8% for households in the 25–29 age range, 6.3% for households in their 30s, 6.5% for households in their 40s, 5.3% for those aged 50–54, and 2.7% for those aged 55–59.

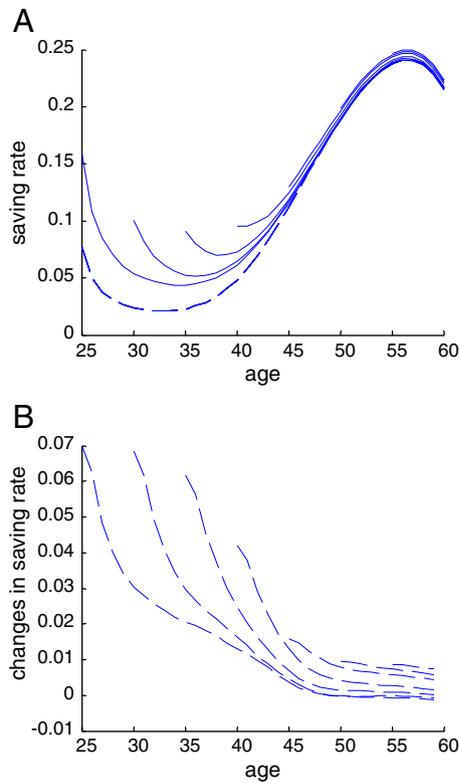


Fig. 4. A. Simulated age profile of saving rates before and after rise in variance of transitory income shock: one-off change. Notes: Dashed line corresponds to the saving behavior when the variance of transitory income shocks is 0.03. Other lines indicate saving behavior when that variance is 0.09 if the change were to occur when the household was at the age where the line begins. B. Jump in saving rate following rise in variance of transitory income shock: one-off change. Notes: Each line corresponds to the jump in the saving rate in panel A after the increase in the variance of transitory income shock for a household at that age.

more muted because those households may already have accumulated a buffer stock of savings. Incorporating an increase in the permanent variance of income would yield an even stronger saving response by households with younger heads.

Fig. 5 is analogous to Fig. 4, but captures the shock to the pension replacement rate. The initial baseline profile in Fig. 5A is the same as in Fig. 4A. The additional lines correspond to simulated age-savings profiles following the decline in the retirement replacement ratio. Fig. 5B plots the change in saving rates relative to that baseline. The change in the replacement ratio induces a substantial increase in savings, particularly for households with older household heads nearing retirement. After the pension reform, households need to save more in order to attain the same level of post-retirement consumption as in the pre-reform scenario. The older the household head, the less time there is to adjust life cycle savings to the lower replacement ratio (i.e., compensate for past savings that were not made because the individual was living in a more favorable pension environment). Hence, while the increase in the saving rate relative to the pre-reform baseline is less than 1 percentage point for a household with a young household head, it is 5 percentage points for those with a household head in his or her mid-forties, and as high as 8 percentage points for households with heads close to the end of their working life.

In practice, the transition from the old to the new system was smoother than the discrete change in our calibration, which therefore overstates the initial jump for the older workers (we discuss this further in Section 6.2). But note that even the young households, which have more time to adjust, will be saving over 5 percentage points more by the time they approach retirement than they would have had under the old pension regime. These long-run estimates are less sensitive to

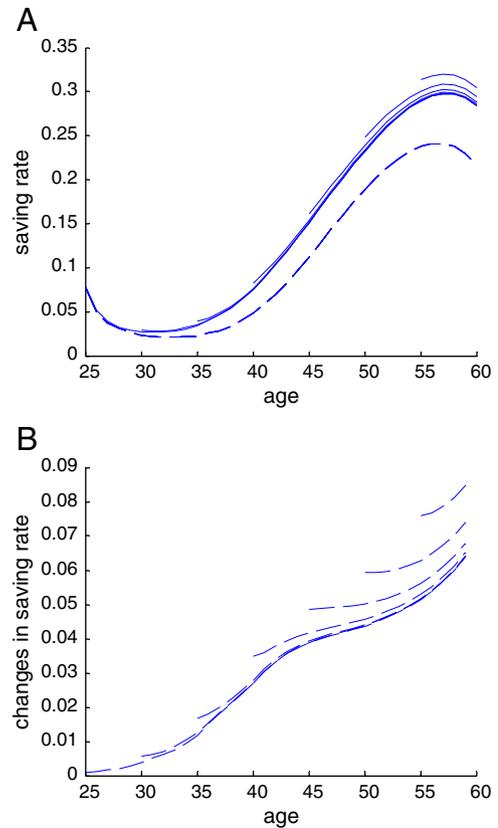


Fig. 5. A. Simulated age profile of saving rates before and after pension reform: one-off change. Notes: Dashed line corresponds to the saving behavior when $\eta = 0.75$. Other lines indicate the saving behavior when $\eta = 0.60$ if the change in pension replacement were to occur when the individual was at that respective age. B. Jump in saving rate following pension reform: one-off change. Notes: Each line corresponds to the jump in the saving rate in panel A after pension reform for a household at that age.

the assumption of an initial discrete change and indicate that, in the long term, households with older household heads will continue to save substantially more than they used to in the past. It is also worth emphasizing that the effect from pension reform could be amplified by the existence of income uncertainty. Uncertain income streams during the household head's working life leads to uncertainty about pension benefits. With smaller anticipated pension benefits, this would lead to higher savings even for households who have many years before retirement.

To jointly evaluate the effects of the rise in transitory income uncertainty as well as the change in the pension replacement rate, in Fig. 6 we show the results when both factors are introduced simultaneously. As expected, saving rates respond more strongly once both shocks are introduced, although the combined result is less than the sum of the two effects from Figs. 4 and 5. The reason for this interaction is that the higher buffer-stock savings accumulated in response to the increase in transitory earnings reduces the need for life-cycle savings later on (and higher life-cycle savings also help protect against temporary shocks to income). There is a marked increase in saving rates at the time of the switch, amounting to about 5 percentage points for households with household heads in their twenties, thirties and forties. Saving rates tend to decline after the initial jump for the younger households (since the main motive for the initial jump is to quickly build up an adequate buffer stock). But for households with heads aged in the mid-forties onwards, saving rates remain more stable (since the retirement motive is already sufficiently strong). And as expected the initial jump (and subsequent saving behavior) is very high for those with heads in their fifties, as those households have less time to adjust to a less generous pension replacement rate.

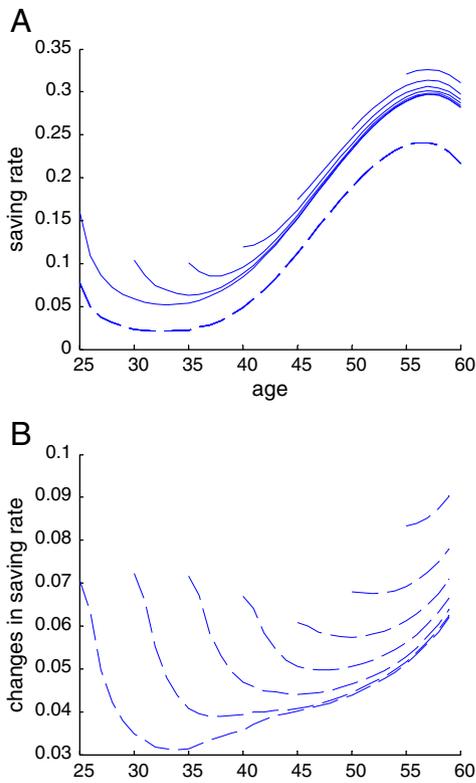


Fig. 6. A. Simulated age profile of saving rates before and after rise in variance of income and pension reform: one-off change. Notes: Dashed line corresponds to the saving behavior when the variance of transitory income is 0.03 and $\eta = 0.75$. Other lines correspond to saving behavior when the variance of transitory income is 0.09 and $\eta = 0.60$ if the change were to occur when the individual was at that respective age. B. Jump in saving rate following rise in variance of income and pension reform: one-off change. Notes: Each line corresponds to the jump in the saving rate in panel A after the shock and pension reform occurs for a household at that age. Changes in saving rates are smoothed by a moving average with 3 neighboring age averages.

The results from Fig. 6 are informative but it is difficult to compare the increase in saving rates from those plots with the increase observed in the cross-sectional data since saving rates in the cross section involve a combination of the initial jumps as well as movements along the curves over time. To facilitate that comparison, Fig. 7 plots the change in cross-sectional saving rates implied by the saving behavior in Fig. 6 at different points in time relative to a discrete increase in uncertainty

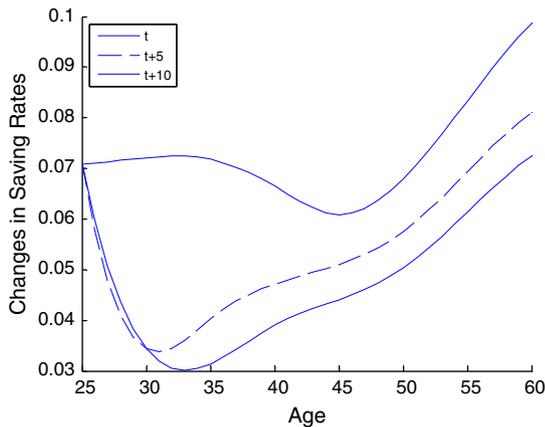


Fig. 7. Projected cross-sectional changes in saving rates after rise in the variance of income and pension reform: one-off change. Notes: Lines correspond to changes in the cross-sectional age profile of savings implied by the profiles in Fig. 6 (households immediately adjusting after a single shock). The implied aggregate increase in the saving rate at t , $t + 5$ and $t + 10$ are 7.6%, 5.2% and 4.8% respectively.

and pension reform. The three lines indicate the change in the cross-section at the time of the change and initial jump (t), as well as at five ($t + 5$) and ten ($t + 10$) year horizons, respectively, after the shifts in income uncertainty and pensions.

Note that even though the envelope of the initial adjustment in savings is non-decreasing for household head aged under 35 in Fig. 6B (consistent with the change in the cross-section at time t), the movement along those lines after the initial adjustment implies a U-shaped pattern for the change in savings in the cross-section at $t + 5$ and $t + 10$. In all plots, the households with the youngest heads save about 7 percentage points more than they used to, while the oldest save 6.5 percentage points more at $t + 5$ and 5.5 percentage points more at $t + 10$. Both the $t + 5$ and $t + 10$ age profiles initially decline rapidly with age, bottoming out at around 3 percentage points for households with heads around age 35 or in their early 30s. In the $t + 10$ cross-section, a household whose head is in his or her early 40s saves 3–4 percentage points more than before the reform, giving the cross-sectional profile an asymmetric U-shaped pattern (with a relatively rapid initial decline followed by a gradual increase in saving rates plotted against age of household head).²⁸

6.2. Gradual reforms and learning

For the sake of simplicity, we have modeled the increase in uncertainty and the pension reform as a composite one-off shock. In practice, the change in uncertainty was gradual. Even major events like the pension reform came to be expected (so workers could respond in anticipation) and the implementation itself was gradual, unlike in our stylized setting. We now consider a scenario where both the increase in the transitory variance and the change in the pension replacement rate take place in two steps. Half of the change (both in the increase in transitory variance and the change in the replacement rate) takes place at time t , with the remaining half taking place at $t + 5$.

Moreover, we also allow for gradual learning by households, with some households learning about the shocks faster than others. This learning is assumed to be independent of household characteristics. One-fifth of the households learn immediately and respond to the initial shock at time t , while the remaining households continue to optimize under the old parameters. At time $t + 1$, another fifth of the households discover the new parameter values and re-optimize and so on until all households have learned the new parameters by time $t + 4$, right before the second change takes place. In this way, between t and $t + 5$, there is always a subpopulation responding to the discrete change in economic environment. This simple learning process implies that the cross-sectional saving rate would shift up gradually, as opposed to a large adjustment at time t followed by a decline starting from $t + 1$ onwards in the baseline version. The same type of learning applies for the second shock at $t + 5$, with households learning in the same sequence. This combination of gradual shocks and learning allows for more realistic dynamics, which more closely tracks the evolution of the age profile of savings in the data.

Fig. 8A and B are analogous to Fig. 6A and B, but plot the results under this gradual reform with learning scenario. These figures have two sets of lines at each age range: one to indicate the initial adjustment in savings, and the other to indicate the adjustment that takes place 5 years later. Note that they provide a snapshot for only one-fifth of the households (those that learn at t and $t + 5$, respectively, about the shocks; we do not plot the behavior of households that learn about the shocks only in later time periods). As in the previous scenario, the

²⁸ In principle, the changes in saving rates for households with heads in the 25–30 age range should be identical at times $t + 5$ and $t + 10$. Fig. 7 shows that, at ages 25 and 30, the simulated changes in savings at $t + 5$ and $t + 10$ are exactly the same. For computational reasons, we solve the model at seven discrete age points (at ages 25, 30, 35 and so on) and then use a spline to smooth the age-saving profile. This smoothing procedure results in a small discrepancy between the lines corresponding to times $t + 5$ and $t + 10$ for households with heads that are 26–29 years old.

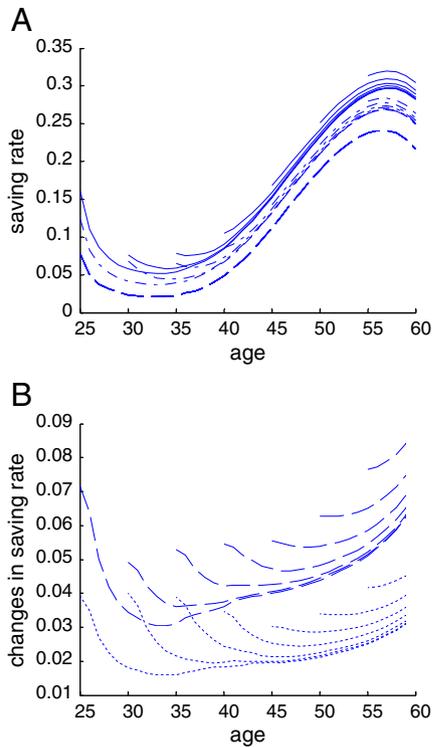


Fig. 8. A. Simulated age profiles of saving rates before and after changes in income risk and pension reform benefits: two-step changes with gradual learning. Notes: Dashed line corresponds to the saving behavior when the variance of transitory income is 0.03 and $\eta = 0.75$. Dash-dotted line corresponds to the saving behavior when the variance of transitory income is 0.06 and $\eta = 0.675$ if the change were to occur at time t when the individual was at that respective age. Solid lines correspond to saving behavior when the variance of transitory income is 0.09 and $\eta = 0.60$ if the change were to occur at time $t + 5$ when the individual was at that respective age. Note that figure provides a snapshot for only one-fifth of households (those that learn at t and $t + 5$). B. Jump in saving rate following changes in income risk and pension reform benefits: two-step changes with gradual learning. Notes: Each line corresponds to the jump in the saving rate in panel A after the shock and pension reform occurs for a household at that age. Dotted lines refer to changes in savings at the first step change at time t . Dashed lines refer to changes in savings at the second step change at time $t + 5$. Changes in saving rates are smoothed by a moving average with 3 neighboring age averages.

combined effect on aggregate savings involves a combination of the initial jumps and movements along the curves over time (with the added complication that, with gradual reforms and learning, a tenth of the households will be discretely adjusting their savings in each of the first ten years).

That combined effect can be more easily observed in Fig. 9, which plots the change in cross-sectional saving at different points in time under this two-step increase in uncertainty and pension reform and gradual learning about those shifts. The four lines indicate the change in the cross-section at the time of the change and initial jump (t), one year after the jump, as well as at one ($t + 1$), five ($t + 5$), and ten ($t + 10$) year horizons, respectively. By time $t + 10$, all changes in income uncertainty and pensions have been learnt by the entire population. The increases in saving rates at time t is 0.75 percentage points, which is one-fifth of the implied initial increases in savings as shown in the dotted line in Fig. 8B. At time $t + 1$, another one-fifth of households learn about changes in uncertainty and pension reforms, leading to a further increase of 1.5 percentage points in the average saving rate. By time $t + 5$, when all households have adapted, the average saving rate rises by 3.3 percentage points and starts becoming U-shaped. By time $t + 10$, when the entire population would have adjusted to the two sets of shocks, the average saving rate rises by 5.6 percentage points compared with the level at time t .

How well do the model predictions match up with the data? The average household saving rate in the Urban Household Survey rose 8.6

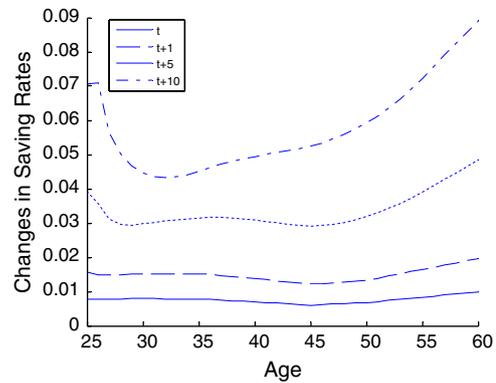


Fig. 9. Projected cross-sectional changes in saving rates changes in income risk and pension reform benefits: two-step changes with gradual learning. Notes: Lines correspond to changes in the cross-sectional age profile of savings implied by the profiles in Fig. 8 and with gradual learning described in Section 6.2. The implied aggregate increases in the saving rate at times t , $t + 1$, $t + 5$ and $t + 10$ are 0.8%, 1.5%, 3.3% and 5.6%, respectively.

percentage points from 1997 to 2007. If we set 1997 as the year t for our calibrations, the 5.6 percentage point increase simulated by our model would explain two-thirds of the actual increase in the savings rate.

Furthermore, the age profile of the saving rate becomes increasingly U-shaped. Households with young and old household heads increase their savings by over 7 percentage points while the increase for households with heads in their 30s is around 4.5 percentage points.

In short, our calibration of a standard buffer-stock/life-cycle model based on parameters taken from our empirical estimation of the shifts in the variances of shocks to household incomes, in combination with estimates of the effects of the 1997 pension reform, can account for a sizable increase in household saving rates and the U-shaped age-savings profile. Chamon and Prasad (2010) trace much of the increase in the saving rates among the young to motives of saving for housing purchases (about 6 percentage points for 25–29 year olds that do not own a home), and among the old to health expenditures (about 6 percentage points for 55–59 year olds). Our calibration exercise suggests that shifts in earnings uncertainty (including the effects of pension reforms) played an important role as well.²⁹ Of course, the numbers we have reported should be viewed as an illustration of the predictive content of a stylized model that has abstracted from other potentially important considerations, rather than a precise calculation. In the next section, we test the robustness of the results to changes in the values of some key parameters.

7. Robustness checks

In this section, we consider a number of variants of the parameters and structure of our model in order to evaluate the robustness of our analytical results. For simplicity, we focus the robustness around the one-off shift discussed in Section 6.1. These results are summarized in Table 4. The first row of that table is our baseline scenario, shown in Fig. 6.

²⁹ Housing motives for saving are not included in the calibration. If included, they would raise the saving rates of the younger individuals, accentuating the U-shaped age-saving profile and bringing it more in conformity with the pattern observed in savings data for Chinese urban households. Lumpy and uncertain health expenditures can still contribute to savings among the elderly (particularly among those that have already retired). But while the inclusion of both effects would further contribute to savings, the combined effect should be smaller than when each is considered in isolation (e.g., a higher buffer-stock accumulated in the aftermath of the pension reform can help older households better cope with health shocks).

Table 4
Robustness checks: simulated rise in savings using various preference parameters.

Jump in saving rates at time of shock for different ages of household head											
	Risk aversion	Discount factor	25	30	35	40	45	50	55	Aggregate increases	
Baseline assumptions	4.5	0.97	8.2%	8.0%	7.8%	7.0%	6.2%	6.8%	8.3%	7.6%	
Larger drop in replacement rate post-reform	4.5	0.97	8.3%	8.4%	8.7%	9.1%	9.4%	10.8%	13.3%	9.3%	
True increase in transitory uncertainty: 30% of estimated increase	4.5	0.97	6.1%	6.2%	6.2%	5.9%	5.6%	6.5%	8.0%	6.3%	
70% of estimated increase	4.5	0.97	9.9%	9.5%	9.1%	8.1%	6.8%	7.2%	8.6%	8.8%	
Baseline changes in replacement rate and uncertainty with alternative preference parameters	2.0	0.97	6.6%	6.8%	6.8%	6.9%	7.3%	8.8%	10.8%	7.3%	
Preferences calibrated based on mean saving rate of 5-year age groups	4.5	0.95	6.9%	8.1%	8.2%	8.2%	7.3%	7.2%	8.7%	7.7%	
	7.5	0.99	2.8%	2.2%	2.5%	3.0%	3.6%	4.4%	5.6%	3.1%	
	8.1	0.97	2.6%	2.0%	2.4%	2.9%	3.4%	4.2%	5.5%	3.0%	

Notes: The first row reports the change in saving behavior when the transitory variance of income increases from 0.03 to 0.09 and the pension replacement ratio declines from 0.75 to 0.60. The column heads refer to age of the household head, so each column shows the impact of those changes for households with household heads of a particular age. The second row evaluates the changes under a scenario proposed by Sin (2005) based on the argument that the drop in the effective replacement rate post-reform is larger (declines to 0.50). The next two rows examine the sensitivity of our assumption about the extent of the true increase in the variance of transitory shocks (as a fraction of the estimated increase in transitory variance; in the baseline, we assume this is 50%). The next set of rows experiments with different values of two key parameters—the risk aversion parameter and the discount factor. The last two rows calibrate preference parameters of risk aversion (while holding the discount factor fixed at 0.99 and 0.97, respectively) so as to match the average age-profile of savings in 1990–1997 using an identity weighting matrix.

7.1. Different shocks to pension replacement rate and variance of transitory income

The second row of Table 4 shows what happens to the saving rate when we assume a larger decline in the pension replacement rate after the reforms, down to a replacement ratio of 50% compared to 60% in our baseline. As expected, that larger decline does not affect the households with youngest heads much on impact (but will eventually affect their savings once they approach retirement), but leads to a significantly larger response among households with older heads.

In the baseline scenario, we considered the effects of an increase in the variance of transitory shocks that was one-half of the empirically estimated increase. This was based on the conservative assumption that half of the increase might represent measurement error. In order to examine the sensitivity of our results to this assumption, we now examine the results when the rise in the transitory income uncertainty is set at 0.3 times and 0.7 times the empirical estimate, respectively (see Table 4, rows 3–4). The results show that the sensitivity of households with younger household heads to changes in the transitory variance is more pronounced. Buffer stock saving motives are strongest for those households as they typically do not yet have savings built up to deal with income volatility.

7.2. Alternative preference parameters

The simulated age profiles of consumption and savings depend critically on the preference parameters as well as on the expected path for income growth. To examine the effects of deviations from the baseline values of the preference parameters, we now simulate the changes in saving using alternative values for the coefficient of risk aversion and the discount factor.

The results in the fifth and sixth rows of Table 4 revert to our baseline assumptions for the increase in uncertainty and pension reform, but show what happens when we vary the risk aversion and discount factor parameters. Across a range of reasonable parameter values, the jump in saving rates is on average broadly comparable to that in our baseline model (although the response for the households with older heads tends to be larger). Lower risk aversion tends to reduce the increase in savings for the young in response to the higher transitory uncertainty. Since a smaller buffer stock is accumulated in the beginning of the life cycle, and at the same time consumers are more willing to substitute away from current consumption towards future consumption, the response to savings can be higher for other age groups for life cycle reasons. In cross-sectional data, the increases in savings would result in a U-shaped age-savings profile after a few years due to the rise in uncertainty and the decline in the replacement ratio.

Next, we calibrate the risk aversion parameter by fitting the simulated age profile of the saving rate to the profile estimated empirically using data from the Urban Household Survey (UHS, which reports income and consumption for different cross-sections of households each year). For this calibration, the empirical cross-section of the saving rates would not be appropriate, since it includes changes due to age, as well as cohort and year effects, and variations in family composition. In order to calibrate the model, we need to estimate the age profile of saving rates while controlling for those other effects.

We construct synthetic cohorts from different cross-sections of the UHS, and regress log income and log consumption on a full set of dummies for age and cohort, and controls for family size (including log of household size and the share of household members in different age groups). We restrict the sample to 1990–1997, since we are trying to calibrate the preference parameters to match the saving behavior prior to the increase in uncertainty. Our estimated age profile is based on the difference between the age effects for log income and the age effects for log consumption.

We then calibrate the risk aversion parameter so as to match the mean saving rate for household heads in seven age groups: 25–29, 30–34, ..., and 55–59, using an identity weighting matrix.³⁰ This matching exercise yields a coefficient of risk aversion of 8.1 and 7.5, when the discount factor is set at 0.97 and 0.99, respectively (see last two rows of Table 4). This high coefficient of risk aversion highlights the challenges of explaining the high saving rates observed in China with a standard buffer-stock life-cycle model of consumption. Given the combination of a generous pension replacement rate, strong expected income growth and relatively low income risk, the only way for the model to capture the high saving rate before 1997 is by setting the risk aversion parameter high enough to reflect very risk-averse consumers. Presumably, more reasonable parameter values could match the observed saving behavior if other saving motives were introduced (e.g., strong bequest and housing motives, and the risk of lumpy health expenditures), which are beyond the scope of this paper. Taking these parameters at face value, the model would still imply an increase in the average saving rate of more than 2.5 percentage points. Note that the increase in savings for households with young household heads is lower than under the baseline scenario, despite the higher risk aversion. Since risk aversion is so high in this scenario, these households already

³⁰ The saving rates used for these seven age groups (from young to old) were: 23.0, 22.7, 21.1, 23.2, 25.9, 24.4 and 22.3%, respectively. This age-saving profile captures the estimated saving behavior for a household with a head aged 25 years in 1997 as he or she ages (not the 1997 cross-section of savings with respect to age). Note that this age-saving profile is not U-shaped as it is based on 1990–1997 data; the U-shaped profile does not appear in the data until the 2000s.

save a lot under the low uncertainty environment, reducing their need to adjust savings once uncertainty rises (which also affects the other age ranges).

Even though it is difficult to explain the actual saving behavior of Chinese households with the standard buffer-stock/life-cycle model, the estimates presented in this section are still useful and informative. These results quantify how far this standard model would go, under reasonable parameter values, towards explaining an increase in savings rates. Our calculations suggest that about half of the increase observed during our sample could be explained by the rise in income uncertainty and pension reform.

7.3. Increase in variance of permanent shocks

Our empirical estimates indicate an increase in the variance of permanent shocks to income but the increase is not statistically significant. Nevertheless, we ran some simulations to evaluate the effects of an increase in the variance of permanent shocks from 0.020 to 0.025 (in addition to the baseline increase in the transitory variance and decline in the pension replacement rate). As expected, an increase in the permanent variance has a much stronger effect on savings than an increase of the same magnitude in the transitory variance. Moreover, while the increase in the transitory variance mainly affects the young (who start with no buffer stock of savings and need to build it up more rapidly), the increase in the permanent variance has a strong effect across all age groups. A permanent shock has implications for retirement earnings since in our simulations the pension is based on the pre-retirement permanent income. Thus, while households with young household heads still need to worry about building their buffer stocks more rapidly following an increase in the permanent variance, households with older heads also need to save more to protect themselves against an adverse permanent shock prior to retirement. For instance, given our baseline parameter values, we find that an increase in the variance of permanent shocks leads to an additional increase of about 4 percentage points in the saving rate of households with younger and older household heads (this is in addition to the increase resulting from the higher variance of transitory shocks and the decline in the pension replacement rate). However, given that our empirical work did not yield robust estimates of a shift in the variance of permanent shocks, we do not pursue this issue further here.³¹

7.4. Changes in expected income growth

Finally, we turn to the issue of how savings may be affected by a slowdown in aggregate income growth. This could happen, for instance, if convergence effects stop propelling growth or labor constraints due to demographic shifts reduce growth. Lower income growth can decrease buffer-stock saving motives (as a lower saving rate is required for the buffer-stock to keep up with permanent income). Lower income growth also reduces the extent to which households postpone their retirement savings towards the end of their life cycle (retirement savings are also affected by how income growth affects the expected retirement period income).

We conduct two experiments. In the first one, expected income growth is set 1 percentage point lower than in the baseline through the entire lifecycle. In the second one, growth is initially the same as in the baseline, but declines by 0.1 percentage points each year relative to the baseline path. This path implies that a 25-year old would have zero income growth by age 58, and we set income growth at zero for the remaining years before retirement. Experiment 1 captures the effects of lower expected income growth for the entire lifecycle while experiment 2 shows the effects of a different trajectory of expected

income growth, which is initially the same as in the baseline but declines much faster.³²

Fig. 10 plots, for different expected income growth scenarios, the age-savings profiles followed by a household with a household head starting off at 25 years of age. Other than the expected growth path, all plots assume the same parameters as in our baseline scenario under the higher uncertainty and lower pension replacement ratio environment. The solid line corresponds to our baseline expected growth path. Slower growth in income causes households with younger heads to save more, with this effect somewhat stronger in experiment 2, when there is a sharper fall-off in the expected income growth rate. To the extent that the young expect slower income growth, they have a stronger incentive to increase their savings for retirement earlier in the life cycle. In our model, that effect is strengthened by the presence of uncertainty since, in addition to retirement, the young also bring forward their buffer-stock savings when expected income growth is lower.

The results suggest that saving rates would be higher for households with young heads under both lower income growth paths, more so when income slows down gradually (dashed line) than when the decline takes place immediately (dotted line). The lower growth path substantially reduces retirement income in our simulation, strengthening retirement saving motives and causing households to save more even in the early stages of the life cycle. The age-savings profile is flatter when the slowdown is gradual than when it takes place immediately. The age-savings profile is steepest under our baseline expected growth path, where postponement of retirement savings is strongest, and savings rates are actually higher in the working years close to retirement age than under both alternative scenarios.

These results suggest that prospects of an eventual slowdown in Chinese growth could further increase household saving rates. Perhaps some of the savings observed among the very young already take into account the prospects of income growth eventually slowing down. Song and Yang (2010) make a similar point based on their results showing a flattening of age-earnings profiles in the UHS data.

8. Conclusion

In this paper, we analyzed a panel of urban Chinese households over the period 1989–2009 and found that the variance of shocks to household income has increased over time and that the increase is mainly accounted for by a rise in the variance of transitory shocks to income. This increase in transitory uncertainty can help explain the rising saving rates among households with younger household heads (who would need larger buffer stocks of savings to handle these shocks). The pension reforms have led to a reduction in pension replacement income relative to average wages for workers retiring after 1997. This cut in the pension replacement ratio can also help explain rising savings rates, particularly for households with older household heads approaching retirement—such households have less time to adjust to the change in pension benefits and must therefore build up an adequate level of savings more quickly.

When we calibrate a standard buffer-stock life-cycle model of consumption with reasonable parameter values, this riskier environment can generate an initial increase of about 7.6 percentage points in the average household saving rate. This one-off shock helps gauge the importance of these channels for savings. To match the time profile of the increase in the saving rate, we also conduct an alternative exercise where we model the shocks as a two-step process, with households gradually learning about the new environment. This setting allows for a more gradual response of the saving rate, which is consistent with

³¹ Detailed results on the effects of an increase in the variance of permanent shocks under different parameter combinations are available from the authors upon request.

³² The income growth path includes the effects of trend income growth as well as age effects on income. Controlling for trend growth, income eventually declines with age, which explains this low value for household income growth despite the strong trend income growth.

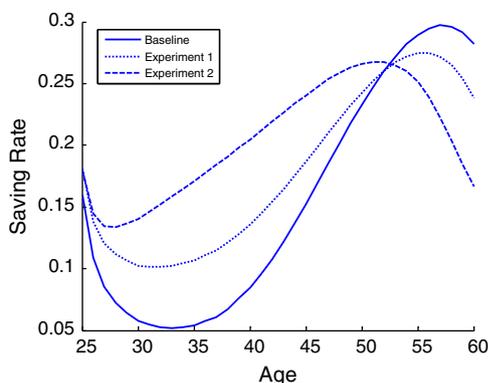


Fig. 10. Simulated age-savings profiles for a household with a young household head under different expected income growth paths: one-off change. Notes: Plot traces expected saving rates over the life cycle for a household with a 25-year old household head under our baseline parameter values after the one-off adjustment for the pension reform and rise in uncertainty described in Section 6.1 (solid line), under a growth path where expected income growth is 1 percentage points lower every year than in our baseline (experiment 1; dotted line) and in a scenario where growth drops by an additional 0.1 percentage points each year relative to the baseline path (with a lower bound of zero income growth; this lower bound is reached at age 58) (experiment 2; dashed line).

the data, and yields an increase in aggregate savings of 5.6% ten years after the initial shock (two-thirds of the increase in aggregate savings observed in the actual data). Moreover, both in the one-off shock and gradual learning settings the increase in savings is concentrated among households with heads at the two ends of the age distribution in our sample. This helps explain the unusual U-shaped age-profile of savings observed in urban China since the late 1990s.

Our calibration involved a number of assumptions for key parameters (e.g., how the pension reform affected the replacement ratio). But we were systematically conservative in our assumptions, erring on the side of downplaying the increase in these risks to income growth. Nevertheless, our calibration of the standard buffer-stock life-cycle consumption model was still capable of explaining half of the observed increase in savings among urban Chinese households, while focusing only on this higher transitory variance of earnings and the 1997 pension reform. These results may be helpful in thinking about policies to rebalance growth in China by boosting private consumption and reducing the reliance on exports and investment to drive growth.

References

- Abowd, John, Card, David, 1989. On the covariance structure of earnings and hours changes. *Econometrica* 57 (2), 411–445.
- Bai, Chong-En, Lu, Jiangyong, Tao, Zhigang, 2006. The multitask theory of state enterprise reform: empirical evidence from China. *American Economic Review* 96 (2), 353–357.
- Baker, Michael, Solon, Gary, 2003. Earnings dynamics and inequality among Canadian men, 1976–1992: evidence from longitudinal income tax records. *Journal of Labor Economics* 21 (2), 289–321.
- Banerjee, Abhijit, Meng, Xin, Qian, Nancy, 2010. Fertility and Savings: Micro Evidence from Family Planning in China. MIT and Yale University Manuscript.
- Bayoumi, Tamim, Tong, Hui, Wei, Shang-jin, 2010. The Chinese corporate savings puzzle: a firm-level cross-country perspective. NBER Working Paper No. 16432.
- Benson, John, Zhu, Ying, 1999. Markets, firms and workers in Chinese state-owned enterprises. *Human Resource Management Journal* 9 (4), 58–74.
- Blanchard, Olivier, Giavazzi, Francesco, 2006. Rebalancing growth in China: a three-handed approach. *China and the World Economy* 14 (4), 1–20.
- Blundell, Richard, Pistaferri, Luigi, Preston, Ian, 2008. Consumption inequality and partial insurance. *American Economic Review* 98 (5), 1887–1921.

- Bound, John, Krueger, Alan, 1994. The extent of measurement error in longitudinal earnings data: do two wrongs make a right? *Journal of Labor Economics* 9, 1–24.
- Brandt, Loren, Hsieh, Chang-Tai, Zhu, Xiaodong, 2008. Growth and structural transformation in China. In: Brandt, Loren, Rawski, Thomas (Eds.), *China's Great Economic Transformation*. Cambridge University Press.
- Carroll, Christopher, 1997. Buffer stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics* 107 (1), 1–56.
- Carroll, Christopher, 2006. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters* 91 (3), 312–320.
- Chamon, Marcos, Prasad, Eswar, 2010. Why are saving rates of urban households in China rising? *American Economic Journal: Macroeconomics* 2 (1), 93–130.
- Comin, Diego, Goshen, Erica L., Rabin, Bess, 2009. Turbulent firms, turbulent wages? *Journal of Monetary Economics* 56 (1), 109–133.
- Ding, Ning, Wang, Yougui, 2008. Household income mobility in China and its decomposition. *China Economic Review* 19 (3), 373–380.
- Dunaway, Steven, Arora, Vivek, 2007. Pension reform in China: the need for a new approach. IMF Working Paper No. 07/109.
- Feng, Jing, He, Lixin, Sato, Hiroshi, 2011. Public pension and household saving: evidence from China. *Journal of Comparative Economics* 39, 470–485.
- Fuchs-Schündeln, Nicola, 2008. The response of household saving to the large shock of German reunification. *American Economic Review* 98 (5), 1798–1828.
- Gang, Fan, Lunati, Maria Rosa, O'Connor, David, 1998. Labour market aspects of state enterprise reform in China. OECD Working Paper, No. 141.
- Gottschalk, Peter, Moffitt, Robert, 1994. The growth of earnings instability in the U.S. labor market. *Brookings Papers on Economic Activity* 2, 217–272.
- Gottschalk, Peter, Moffitt, Robert, 2009. The rising instability of U.S. earnings. *Journal of Economic Perspectives* 23 (4), 3–24.
- Gourinchas, Pierre-Olivier, Parker, Jonathan A., 2002. Consumption over the life cycle. *Econometrica* 70 (1), 47–89.
- Groves, Theodore, Hong, Yongmiao, McMillan, John, Naughton, Barry, 1994. Autonomy and incentives in Chinese state enterprises. *The Quarterly Journal of Economics* 109 (1), 183–209.
- Herd, Richard, Hu, Hu-Wei, Koen, Vincent, 2010. Providing greater old-age security in China. OECD Economics Department Working Paper No. 750.
- Horioka, Charles Yuji, Wan, Junmin, 2007. The determinants of household saving in China: a dynamic panel analysis of provincial data. *Journal of Money, Credit, and Banking* 39 (8), 2077–2096.
- Kaplan, Greg, Violante, Gianluca, 2010. How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics* 2 (4), 53–87.
- Kraay, Aart, 2000. Household saving in China. *World Bank Economic Review* 14 (3), 545–570.
- Kuijs, Louis, 2006. How will China's saving-investment balance evolve? World Bank Policy Research Working Paper No. 3958.
- Li, Weiye, Putterman, Louis, 2008. Reforming China's SOEs: an overview. *Comparative Economic Studies* 50 (3), 353–380.
- Li, Hongbin, Zhu, Yi, 2006. Income, income inequality, and health: evidence from China. *Journal of Comparative Economics* 34 (4), 668–693.
- Lillard, Lee A., Weiss, Yoram, 1979. Components of variation in panel earnings data: American scientists, 1960–70. *Econometrica* 47 (2), 437–454.
- Lin, Justin Yifu, Cai, Fang, Li, Zhou, 1998. Competition, policy burdens, and state-owned enterprise reform. *American Economic Review* 88 (2), 422–427.
- MacCurdy, Thomas E., 1982. The use of time series processes to model the error structure of earnings in longitudinal data analysis. *Journal of Econometrics* 18 (1), 83–114.
- Meghir, Costas, Pistaferri, Luigi, 2004. Income variance dynamics and heterogeneity. *Econometrica* 72 (1), 1–32.
- Meng, Xin, 2003. Unemployment, consumption smoothing, and precautionary saving in urban China. *Journal of Comparative Economics* 31 (3), 465–485.
- Modigliani, Franco, Cao, Shi Larry, 2004. The Chinese saving puzzle and the life cycle hypothesis. *Journal of Economic Literature* 42, 145–170.
- Moffitt, Robert, Gottschalk, Peter, 1995. "Trends in the Transitory Variance of Male Earnings in the U.S., 1970–1987," Mimeo, Johns Hopkins University.
- Pischke, Jorn-Steffen, 1995. Measurement error and earnings dynamics: some estimates from the PSID validation study. *Journal of Business and Economic Statistics* 13 (3), 305–314.
- Prasad, Eswar S., 2011. Rebalancing growth in Asia. *International Finance* 14 (1), 27–66.
- Qian, Yingyi, 1998. Urban and rural household saving in China. IMF Staff Papers 35 (4), 592–627.
- Sin, Yvonne, 2005. Pension liabilities and reform options for old age insurance. World Bank Working Paper No. 2005-1.
- Song, Michael, Yang, Dennis, 2010. Life Cycle Earnings and the Household Saving Puzzle in a Fast-growing Economy. Chinese University of Hong Kong (Manuscript).
- Topel, Robert H., Ward, Michael P., 1992. Job mobility and the careers of young men. *Quarterly Journal of Economics* 107 (2), 439–479.
- Violante, Gianluca, 2002. Technological acceleration, skill transferability, and the rise in residual inequality. *Quarterly Journal of Economics* 117 (1), 297–338.
- Wei, Shang-jin, Zhang, Xiaobo, 2011. The competitive saving motive: evidence from rising sex ratios and savings rates in China. *Journal of Political Economy* 119 (3), 511–564.