



IMF Working Paper

Income Uncertainty and Household Savings in China

Marcos Chamon, Kai Liu, and Eswar Prasad

IMF Working Paper
Research Department

Income Uncertainty and Household Savings in China

Prepared by Marcos Chamon, Kai Liu, and Eswar Prasad*

Authorized for distribution by Atish R. Ghosh
December 2010

This Working Paper should not be reported as representing the views of the IMF.

The views expressed in this Working Paper are those of the author(s) and do not necessarily represent those of the IMF or IMF policy. Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate.

Abstract

China's household saving rate has increased markedly since the mid-1990s and the age-savings profile has become U-shaped. We find that rising income uncertainty and pension reforms help explain both of these phenomena. Using a panel of Chinese households covering the period 1989-2006, we document that strong average income growth has been accompanied by a substantial increase in income uncertainty. Interestingly, the permanent variance of household income remains stable while it is the transitory variance that rises sharply. A calibration of a buffer-stock savings model indicates that rising savings rates among younger households are consistent with rising income uncertainty and higher saving rates among older households are consistent with a decline in the pension replacement ratio for those retiring after 1997. We conclude that rising income uncertainty and pension reforms can account for over half of the increase in the urban household savings rate in China since the mid-1990s as well as the U-shaped age-profile of savings.

JEL Classification Numbers: D91, J3, E21

Keywords: China, household savings, income uncertainty, pension reforms, buffer-stock savings.

Author's E-Mail Address: MChamon@imf.org; klu@jhu.edu; eswar.prasad@cornell.edu.

* Chamon: Research Department, IMF, mchamon@imf.org; Liu: Johns Hopkins University, klu@jhu.edu; Prasad: Cornell University, Brookings Institution and NBER, eswar.prasad@cornell.edu. We are grateful to Loren Brandt, Chris Carroll, Mark Dorfman, Robert Moffitt, Damiano Sandri, 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.

Contents	Page
Abstract.....	2
I. Introduction.....	3
II. Dataset.....	5
III. A Decomposition of Permanent and Transitory Shocks to Labor Income.....	7
IV. Earnings Decomposition Results.....	10
V. Implications of the Shifts in Labor Income Variance for Precautionary Savings.....	13
A. Stylized Facts.....	13
B. The Model and Calibration.....	14
C. Calibration Results.....	16
D. Robustness Checks.....	18
Preference Parameters.....	18
Expected Income Growth.....	20
vi. Conclusion.....	21
Tables	
Table 1. Summary Statistics.....	25
Table 2. Estimated Variance of Permanent and Transitory Shocks to Urban Household Income.....	26
Table 3. Robustness Checks: Simulated Rise in Savings Using Various Preference Parameters.....	27
Figures	
Figure 1. Household Saving Rates.....	28
Figure 2. Gross Saving Rates by Sector.....	28
Figure 3. Urban Household Saving Rates by Age of Household Head.....	29
Figure 4. Labor Market Turnover: Annualized Transition Probabilities.....	29
Figure 5A. Estimated Age Profile of Saving Rates Before and After Rise in Variance of.....	30
Figure 5B. Jump in Saving Rate Following Rise in Variance of.....	30
Figure 6A. Estimated Age Profile of Saving Rates Before and After Pension Reform.....	31
Figure 6B. Jump in Saving Rates Following Pension Reform.....	31
Figure 7A. Estimated Age Profile of Saving Rates Before and After Rise in Transitory.....	32
Figure 7B. Jump in Saving Rates Following Rise in Transitory Variance of Income and.....	32
Figure 8. Projected Cross-Sectional Change in Saving Rates After Rise in the.....	33
Figure 9. Estimated Age Profile of Saving Rates for a Household with 25 Year Old.....	34
References	
References.....	22

I. INTRODUCTION

The Chinese economy has been undergoing a marked transformation in recent decades—from a closed to an open economy, from an agricultural to an industrial economy, and from planned to a more market-oriented economy. This set of transformations has resulted in a rapid growth in average household incomes but has also increased uncertainty as the economy undergoes massive structural shifts. This process 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 shifts on the degree of income uncertainty at the household level and to analyze the implications for household saving rates.

The motivation for this research is that the average saving rate for urban households in China has increased from 15 percent in the early 1990s to over 30 percent in recent years (Figure 1).¹ Households are the main contributors to China's large national savings (Figure 2). This pattern of a rising household saving rate at a time of high income growth seems inconsistent with a certainty-equivalent life-cycle hypothesis model, 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 been a particularly pronounced increase in saving rates among households with younger and older household heads (Figure 3; Chamon and Prasad, 2010). 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 saving rates as well as the dramatic shift in the age-saving profile.

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 urban labor income using a sample of urban households tracked by the China Health and Nutrition Survey (CHNS). 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, Pistaferri and Preston (2008).² We find strong trend growth in both the mean and the variance of total household income. Interestingly, the variance of permanent shocks to household income has remained relatively stable, while the variance of transitory shocks trends upwards. This result is in line with a large literature on how

¹ Other recent studies analyzing the determinants of household savings in China include those using aggregate data (e.g., Modigliani and Cao, 2004; Kuijs, 2006), provincial-level data (e.g., Qian, 1998; Kraay, 2000; Horioka and Wan, 2007; Wei and Zhang, 2009) and micro data at the household or individual levels (e.g., Song and Yang, 2010). Banerjee, Meng and Qian (2010) use a single-year cross-sectional survey to examine the effect of fertility on household savings.

² Other papers in this literature include 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.

technological and sectoral shifts and the associated labor reallocation can generate higher transitory uncertainty even though some of these shifts themselves are permanent in nature. 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 transitory variance of 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. But after that initial adjustment, the response in saving rates gradually 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 pension reform, urban workers received pensions through their employers—predominantly state-owned enterprises. These pensions had a replacement ratio of about 75-80 percent relative to the average wage (Sin, 2005; Arora and Dunaway, 2007). Workers retiring after 1997 are covered under the reformed system. They receive a social pension corresponding to 20 percent of the average local wages, 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 percent of the average wage.³ In our calibration exercise, we show that a decline in the replacement rate from 75 percent to 60 percent 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 to account for the change in pension regime (the initial increase is about one percentage points for households with heads in their thirties). But even the youngest cohorts will 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

³ The social pension is financed by employer contributions of 17 percent of wages. The individual accounts are financed by employer and employee contributions of 3 percent and 8 percent of wages, respectively. Sin (2005) and Arora and Dunaway (2007) provide more details on various aspects of the pension reform, including the gaps in coverage under the new system. Herd, Hu and Koen (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.

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, 2009). 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 2.5-3.5 percentage points more for those in their late thirties-early forties. Our results are robust to alternative parameterizations of the model and conservative assumptions about the rise in transitory income uncertainty.

II. DATASET

We use data from the China Health and Nutrition Survey (CHNS).⁴ 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. This survey was conducted in 1989, 1991, 1993, 1997, 2000, 2004 and 2006.

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 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. 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 urban residents, between the ages of 25 to 59, not a student, and for whom we have complete information on age and education. We drop the small number of urban households reporting income from farming and raising livestock. We include households in every year in which they appear in the data and satisfy these requirements. To limit the effects of extreme observations, we run simple Mincerian earnings regression for each survey year (discussed in detail in the next section), sort the sample in each survey year by the growth in

⁴ 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>

income residuals, and trim out the households in the top and bottom 1 percent of that distribution. Our final sample is an unbalanced panel consisting of 1689 households.⁵

Table 1 shows the number of observations in each year and also presents some summary statistics for the analysis sample. From 1989 to 2006, real mean annual household income almost triples, from 12830 to 32040 RMB at constant 2006 prices.⁶ 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—which refers to government units, state-owned enterprises, and large collective enterprise (with a provincial or city government as the principal owner)—still plays an important role in the Chinese economy.⁷ In our sample, the proportion of workers employed in the SOCE sector falls from 81 percent in 1989 to 64 percent in 2006. Indeed, reflecting the pick-up in the pace of the shift towards a market-oriented economy in recent years, the probability of transition from SOCE to private sector employment rises from 5.5 percent between the years 1989 and 1993 to 14 percent between the years 2000 and 2004. This increase could be due to employment shifts from the SOCE sector to the private sector and also because SOCEs are being restructured and turned into privately-owned firms. We cannot evaluate the relative importance of these two factors as the CHNS does not track job changes. Nevertheless, this is not a big concern from our perspective as either explanation could be consistent with changes in income uncertainty.

There are some concerns about the quality of the income information in the first wave of the CHNS. The administrators of the CHNS had admitted problems in the 1989 wave of the survey, although they note subsequently that “those problems have been corrected and the CHNS 1989 data are now of high quality”. In our analysis sample, the standard deviations of individual earnings and household income are both much larger in 1989 compared with the next couple of waves of the panel. This runs counter to the literature documenting the rise in income inequality in China since the 1980s. Benjamin, Brandt, Loren and Fan (2003) raise similar concerns about

⁵ Gottschalk and Moffitt (2009) discuss the need for trimming. The sample is unbalanced because attrition, new respondents introduced 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.

⁶ In our sample of urban households, labor earnings comprise 97 percent of total individual income on average, so these two measures are approximately equivalent.

⁷ 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, Cai and Li (1998), Bai, Lu and Tao (2006) and Li and Putterman (2008). Brandt, Hsieh and Zhu (2008) analyze the effects of the reallocation of labor from the state sector to the non-state sector.

the wage data collected in the first wave of CHNS. Our baseline results are based on the full panel including 1989, but we will later check the robustness of the results by dropping 1989 from the sample.

III. A DECOMPOSITION OF PERMANENT AND TRANSITORY SHOCKS TO LABOR INCOME

In this section, we describe the methodology we use to decompose the variance in labor income into the components attributable to permanent and temporary shocks. We focus on household labor income, which is more relevant for household consumption and saving decisions, rather than individual labor earnings. Following the literature modeling earning dynamics, we first run Mincerian income regressions that allow us to control for year by year cross-sectional income variation attributable to economy-wide shifts in the returns to observed household characteristics. We regress log income on four region dummies (East, Northeast, Midwest and West), age and education level of the household head (for the latter, we use three education dummies--middle school or less, high school, some college), dummies for the number of income earners in the household, household size, and dummies for the employment and marriage status of the household head. This regression is run separately for each year.⁸

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:

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

where y_{iat} is the log earning residuals for household i aged a in year t from the Mincerian regression, u_{iat} is the permanent component, and v_{iat} is the transitory component. 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.⁹ We assume:

⁸ To conserve space, we do not report these 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 (including adding a dummy for SOCE employment) 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.

⁹ 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): -0.101 (0.011), 0.007 (0.014), -0.000

(continued...)

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

In other words, the variances of permanent and transitory shocks change by year but do not depend on age. This in effect amounts to averaging over households with different ages (or in different cohorts) in each year.¹⁰ Later, we will examine how these variances differ across age groups. From here on, the subscript a will be dropped. The parameters to estimate are: $\sigma_{\xi t}^2$ and $\sigma_{\omega t}^2$ for each survey wave: $t = \{1989, 1991, 1993, 1997, 2000, 2004, 2006\}$.

Suppose we observe household income in consecutive years. Identification hinges on the variance and covariance structure of one-year changes in income (see, e.g., Meghir and Pistaferri, 2004; Blundell, Pistaferri and Preston, 2008):

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

Thus, the one-period lagged autocovariance of income changes identifies the variance of the transitory shock. 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 $\sigma_{u_0}^2$, 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 $\sigma_{u_0}^2$ is likely to be nonstationary.

(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 the transitory shocks are either i.i.d or follow an MA(1) process. The latter is consistent with much of the literature (Abowd and Card, 1989; Meghir and Pistaferri, 2004; Blundell, Pistaferri and Preston, 2008). 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.

¹⁰ 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 uneven spacing of the CHNS waves complicates the analysis since we need to use n-year rather one-year income changes:

$$\text{cov}(\Delta 93 - 91, \Delta 91 - 89) = -\sigma_{\xi 91}^2$$

$$\text{cov}(\Delta 97 - 93, \Delta 93 - 91) = -\sigma_{\xi 93}^2$$

$$\text{cov}(\Delta 00 - 97, \Delta 97 - 93) = -\sigma_{\xi 97}^2$$

$$\text{cov}(\Delta 04 - 00, \Delta 00 - 97) = -\sigma_{\xi 00}^2$$

$$\text{cov}(\Delta 06 - 04, \Delta 04 - 00) = -\sigma_{\xi 04}^2$$

$$\text{var}(\Delta 93 - 91) = \sigma_{\omega 93}^2 + \sigma_{\omega 92}^2 + \sigma_{\xi 93}^2 + \sigma_{\xi 91}^2$$

$$= \sigma_{\omega 93}^2 + \sigma_{\omega 92}^2 - \text{cov}(\Delta 97 - 93, \Delta 93 - 91) - \text{cov}(\Delta 93 - 91, \Delta 91 - 89)$$

...

$$\text{var}(\Delta 04 - 00) = \sigma_{\omega 04}^2 + \sigma_{\omega 03}^2 + \sigma_{\omega 02}^2 + \sigma_{\omega 01}^2 + \sigma_{\xi 04}^2 + \sigma_{\xi 00}^2$$

$$\text{var}(\Delta 06 - 04) = \sigma_{\omega 06}^2 + \sigma_{\omega 05}^2 + \sigma_{\xi 06}^2 + \sigma_{\xi 04}^2$$

We are able to identify five years of the transitory income risk, all except 2006 and 1989. We do not make any assumption about the transitory variances in those two years and, hence, we are able to identify four permanent variances:

$$\sigma_{\omega 93}^2 + \sigma_{\omega 92}^2$$

$$\sigma_{\omega 97}^2 + \sigma_{\omega 96}^2 + \sigma_{\omega 95}^2 + \sigma_{\omega 94}^2$$

$$\sigma_{\omega 00}^2 + \sigma_{\omega 99}^2 + \sigma_{\omega 98}^2$$

$$\sigma_{\omega 04}^2 + \sigma_{\omega 03}^2 + \sigma_{\omega 02}^2 + \sigma_{\omega 01}^2$$

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, transitory shocks may follow an MA(1) process, implying that the identified transitory variance also includes the transitory shocks from the previous year. However, it's worth noting that 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.¹¹ To see this:

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

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 the estimated MA(1) parameter for the transitory shock is small (between -0.1 to -0.2; see, e.g., Blundell, Pistaferri and Preston, 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 percent of the variance of the growth of income. 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.¹²

If we drop the 1989 wave of the CHNS (due to potential data concerns described in the previous section), we are no longer able to estimate the variance of the transitory shock in 1991 and the variance of the permanent shock in 1993. The model is still just identified. The rest of the estimated variances remain the same, since identification of these parameters does not hinge on the income information from 1989 (see the formulas above; for example, identification of the transitory uncertainty in 1993 only requires income data from 1991, 1993 and 1997). While we present our estimates for the whole sample, it is important to exercise caution when interpreting the transitory variance estimates in 1991 and the permanent variance estimates in 1993, whose identification depends on income data from 1989.

IV. EARNINGS DECOMPOSITION RESULTS

Table 2 reports, in panels A and B respectively, estimates for the variances of the permanent and

¹¹ However, without additional assumptions, it is not possible to identify the MA(1) process given the data availability in our sample.

¹² 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.

transitory shocks to household income and earnings over time. Standard errors are computed using a block bootstrap procedure. 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. The same is true if we turn to a sample of households whose head worked in the SOCE sector when the household entered the panel.¹³ When we split the sample by age (head above or below forty years of age when the household enters the panel), the results remain broadly similar, suggesting no clear trend (there is a large dip in 2000 for the older subsample, but one must bear in mind the wide standard errors). 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.

For the results reported in the two penultimate columns, we revert to the full sample but consider alternative definitions of income. The third column from the end shows results for an income measure excluding transfers and subsidies while the results in the penultimate column are based solely on labor earnings. None of the income measures yields evidence of a clear trend in the variance of permanent shocks.

In Panel B, we present estimates of the variance of transitory shocks to household income and earnings. For the full sample, there is a clear trend of a rising variance, with the point estimates steadily rising from 0.040 in 1991 to 0.162 by 2004. The results are similar in the SOCE sample. The subsamples where we divide observations by age 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 (third column from the end).¹⁴ 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. When we further exclude other household income (penultimate column), the rising trend in transitory uncertainty is similar to other measures, except for a partial reversal of the

¹³ 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).

¹⁴ Transfers include both private and public transfers. Subsidies constitute firm-level nonwage compensation to the worker, and include subsidies on gas, food, education and housing as well as allowances for children. 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 percent in the 1991 wave, but steadily declines to about 5 percent in the 2006 wave.

increase in the last survey year. That result reflects the reduced volatility of “other income.”¹⁵

As discussed in the previous section, there are some concerns about the data in 1989. To evaluate the potential impact of data issues in 1989, we re-estimated the model dropping data for that year, which limits identification of the variance of permanent shocks to 1997 and later and the variance of temporary shocks to 1993 and later. The point estimates reported in the last column of Table 2 confirm the trend of rising transitory income uncertainty.

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.

Figure 4 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 uncertainty began rising.¹⁶ The transition from employment in the SOCE sector to employment in the non-SOCE sector is also high starting in the mid-1990s. 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 an increase in the 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 and more openness to foreign trade (see Groves, Hong, McMillan and Naughton, 1994; Gang, Lunati and O’Connor, 1998; Benson and Zhu, 1999).

Comin, Groshen, and Rabin (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-

¹⁵ In the survey, other income refers to income from leased land (only for 1989), rent from non-land assets, rent from lodgers or boarders, and other unspecified sources of income.

¹⁶ In the CHNS, the question concerning employment status is: “Are you working at present?” which is then followed up by a question regarding the reasons for not working. The question is consistent across survey waves although it is only indicative of the employment status for a worker in a particular year.

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.

V. IMPLICATIONS OF THE SHIFTS IN LABOR 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, 2009; 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).¹⁷ 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.

A. Stylized Facts

To motivate this exercise, we turn again to Figure 3, 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).¹⁸ In the early 1990s, the age-saving profile in China was fairly typical of those in other economies, with saving rates increasing with age and then dropping off after retirement. Over time, savings rates have increased across the board. But the increase is particularly pronounced for households with

¹⁷ 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.

¹⁸ The sample covers the following provinces: Anhui, Beijing, Chongqin, Ganshu, Guangdong, Hubei, Jiangsu, Liaoning, Shanxi and Sichuan. Only three of these overlap with the CHNS sample.

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 saving 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.

B. The 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:

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

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

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 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:

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

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

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 percent per annum, which

matches the average real interest rate in China over the period 1989-2006.¹⁹ The coefficient of relative risk aversion γ is 4.5. We assume that people 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. There are no bequests (for an individual who dies prior to age 85 with a positive level of assets, those assets are “lost”).²⁰

Permanent income in the retirement period is initially set such that $\eta=75$ percent of pre-retirement permanent income, which is 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 $\eta=60$.²¹

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 age. We calculate the marginal effect of age on household income at each age.²² The predicted annual income growth is about 7 percent for the young, then ranges between 6 and 7 percent throughout most of the remaining work life before gradually declining to 2 percent close to retirement age.

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 0.04 in the baseline case up to 0.08. 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.

¹⁹ The real interest rate is based on the nominal interest rate on one-year bank deposits deflated by the annual CPI inflation rate.

²⁰ 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.

²¹ 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 percent of the average wage by 2010, 2020 and 2030 respectively. Thus, our assumption for the decline in the replacement rate is a fairly conservative one, particularly for the younger workers.

²² 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.

C. Calibration Results

Figure 5A plots the simulated age profile of the saving rate. We construct the age-saving profile by simulating the model for 5,000 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 savings rate.

The additional lines in this figure correspond to the age-saving 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 Figure 5B, which plots the change in the saving 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 more muted because those households may already have accumulated a buffer stock of savings.

Figure 6 is analogous to Figure 5, but captures the shock to the pension replacement rate. The initial baseline profile in Figure 6A is the same as in Figure 5A. The additional lines correspond to simulated age-saving profiles following the decline in the retirement replacement ratio. Figure 6B plots the change in saving rates relative to that baseline. The change in the replacement ratio induces a substantial increase in savings, particularly for the 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). As a result, 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 (in his or her early thirties), 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 as a result overstates the initial jump for the older workers. But note that even the young households, which have plenty of 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 much less sensitive to 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.

To jointly evaluate the effects of the rise in transitory income uncertainty as well as the change in the pension replacement rate, in Figure 7 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 Figures 5 and 6. 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.

The results from Figure 7 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, Figure 8 plots the change in cross-sectional saving rates implied by the saving behavior in Figure 7 at different points in time relative to a discrete increase in uncertainty 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 $t+5$ and $t+10$ years, which represent 5 and 10 year horizons, respectively, after the shifts in income uncertainty and pensions. Note that even though the envelope of the initial adjustment in savings increases with age in Figure 5B (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 5 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 2.5 percentage points for households with heads around age thirty or in their early thirties. In the $t+10$ cross-section, a household whose head is in his or her early forties 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 savings rates plotted against age of household head).

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 labor earnings, in combination with estimates of the effects of the 1997 pension reform, can account for a sizable increase in saving behavior. 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.²³

D. Robustness Checks

Preference Parameters

The simulated age profiles of consumption and savings depend critically on the preference parameters--the coefficient of risk aversion and the discount factor--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 various combinations of these parameters.

Table 3 shows, for different values of key parameters, the simulated initial rise in saving rate for different ages when there is both a rise in transitory uncertainty and a decline in the pension replacement ratio as specified in the previous section. The first row is our baseline scenario, shown in Figure 7. The second row assumes a larger decline in the pension replacement rate after the reforms, down to a replacement ratio of 50 percent compared to 60 percent 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 for the ones with older heads.

The results in the remaining rows of Table 3 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

²³ 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).

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-saving profile after a few years due to the rise in uncertainty and the decline in the replacement ratio.

In the penultimate row of Table 3, 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 the saving rates while controlling for those other effects.

We construct synthetic cohorts from the 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). But 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.²⁴ 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. 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

²⁴ 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 percent, 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.

households with young households 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 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 saving 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.

Expected Income Growth

Finally, we turn to the issue of how savings may be affected by an eventual slowdown in aggregate income growth in China. This could happen, for instance, if convergence effects stop propelling growth in China or labor constraints due to demographic shifts cool off 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).

Figure 9 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. The dotted line corresponds to a growth path that is 1 percentage point lower than in our baseline from the present onwards. The dashed line is based on a growth path where the expected income growth flattens out at a rate of 0.1 percentage point every year relative to its value in our baseline path, but never declines below the lowest value in that path (1.26 percent).²⁵

The plots indicate that saving rates would be higher for households with young heads under both lower income growth paths, more so when income slows down gradually (dotted line) than when the decline takes place immediately (dashed 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

²⁵ 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 income growth despite the strong trend income growth.

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 saving 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.

VI. CONCLUSION

In this paper, we analyzed a panel of urban Chinese households to show 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 saving 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 is capable of generating an initial increase of about four and a half percentage points in the average household saving rate. Saving rates remain significantly higher after that initial adjustment, with an average increase across age groups of four and a half percentage points after 5 years. Moreover, this increase is concentrated among households with household heads at the two ends of the age distribution in our sample. This also 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, and David Card, 1989, "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, Vol. 57, No. 2, pp. 411–45.
- Bai, Chong-En, Jiangyong Lu, and Zhigang Tao, 2006, "The Multitask Theory of State Enterprise Reform: Empirical Evidence from China," *American Economic Review*, Vol. 96, No. 2, pp. 353–57.
- Banerjee, Abhijit, Xin Meng, and Nancy Qian, 2010, "Fertility and Savings: Micro Evidence from Family Planning in China," Manuscript, MIT and Yale University.
- Baker, Michael, and Gary Solon, 2003, "Earnings Dynamics and Inequality Among Canadian Men, 1976-1992: Evidence from Longitudinal Income Tax Records," *Journal of Labor Economics*, Vol. 21, No. 2, pp. 289-321.
- Benson, John, and Ying Zhu, 1999, "Markets, Firms and Workers In Chinese State-Owned Enterprises," *Human Resource Management Journal*, Vol. 9, No. 4, pp. 58–74.
- Blanchard, Olivier, and Francesco Giavazzi, 2006, "Rebalancing Growth in China: A Three-handed Approach," *China and the World Economy*, Vol. 14, No. 4, pp. 1–20.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston, 2008, "Consumption Inequality and Partial Insurance," *American Economic Review*, Vol. 98, No. 5, pp. 1887–921.
- Bound, John, and Alan Krueger, 1994, "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?" *Journal of Labor Economics*, Vol. 9, pp. 1–24.
- Brandt, Loren, Chang-Tai Hsieh, and Xiaodong Zhu, "Growth and Structural Transformation in China," in Loren Brandt and Thomas Rawski, eds., 2008, *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*, Vol. 107, No. 1, pp. 1–56.
- Carroll, Christopher, 2006, "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, Vol. 91, No. 3, pp. 312–20.
- Chamon, Marcos, and Eswar Prasad, 2010, "Why are Saving Rates of Urban Households in China Rising?" *American Economic Journal: Macroeconomics*, Vol. 2, No. 1, pp. 93–130.
- Comin, Diego, Erica L. Groshen and Bess Rabin, 2009, "Turbulent Firms, Turbulent Wages?" *Journal of Monetary Economics*, Vol. 56, No. 1, pp. 109–33.
- Dunaway, Steven, and Vivek Arora, 2007, "Pension Reform in China: The Need for a New

- Approach,” IMF Working Paper No. 07/109.
- Dwayne, Benjamin, Loren Brandt, and Jia-Zhueng Fan, 2003, “Ceaseless Toil? Health and Labour Supply of the Elderly in Rural China,” Working paper, University of Toronto.
- Feldstein, Martin, 1999, “Social Security Pension Reform in China,” *China Economic Review*, Volume 10, No.2, pp. 99-107.
- Feng, Jing, Lixin He, and Hiroshi Sato, 2009, “Public Pension and Household Saving: Evidence from China,” Bank of Finland BOFIT Discussion Paper 02/2009.
- Fuchs-Schündeln, Nicola, 2008, “The Response of Household Saving to the Large Shock of German Reunification,” *American Economic Review*, Vol. 98, No. 5, pp. 1798–828.
- Gang, Fan, Maria Rosa Lunati, and David O’Connor, 1998, “Labour Market Aspects of State Enterprise Reform in China,” *OECD Working Paper*, No. 141.
- Gottschalk, Peter, and Robert Moffitt, 1994, “The Growth of Earnings Instability in the U.S. Labor Market,” *Brookings Papers on Economic Activity* 2, pp. 217–72.
- Gottschalk, Peter, and Robert Moffitt, 2009, “The Rising Instability of U.S. Earnings,” *Journal of Economic Perspectives*, Vol. 23, No. 4, pp. 3–24.
- Gourinchas, Pierre-Olivier, and Jonathan A. Parker, 2002, “Consumption Over the Life Cycle,” *Econometrica*, Vol. 70, No. 1, pp. 47–89.
- Groves, Theodore, Yongmiao Hong, John McMillan and Barry Naughton, 1994, “Autonomy and Incentives in Chinese State Enterprises,” *The Quarterly Journal of Economics*, Vol. 109, No. 1, pp. 183–209.
- Herd, Richard, Hu-Wei Hu and Vincent Koen, 2010, “Providing Greater Old-Age Security in China”, *OECD Economics Department Working Papers*, No. 750.
- Horioka, Charles Yuji, and Junmin Wan, 2007, “The Determinants of Household Saving in China: A Dynamic Panel Analysis of Provincial Data,” *Journal of Money, Credit, and Banking*, Vol. 39, No. 8, pp. 2077–96.
- Kaplan, Greg, and Gianluca Violante, 2010, “How Much Consumption Insurance Beyond Self-Insurance?” *American Economic Journal: Macroeconomics*, Vol. 2, No. 4, pp. 53–87.
- Kraay, Aart, 2000, “Household Saving in China,” *World Bank Economic Review*, Vol. 14, No.3, pp. 545–70.
- Kuijs, Louis, 2006, “How Will China’s Saving-Investment Balance Evolve?” World Bank Policy Research Working Paper, No. 3958.

- Li, Weiye, and Louis Putterman, 2008, "Reforming China's SOEs: An Overview," *Comparative Economic Studies*, Vol. 50, No. 3, pp. 353–80.
- Lillard, Lee A., and Yoram Weiss, 1979, "Components of Variation in Panel Earnings Data: American Scientists, 1960–70," *Econometrica*, Vol. 47, No. 2, pp. 437–54.
- Lin, Justin Yifu, Fang Cai, and Zhou Li, 1998, "Competition, Policy Burdens, and State-Owned Enterprise Reform," *American Economic Review*, Vol. 88, No. 2, pp. 422–27.
- MaCurdy, Thomas E., 1982, "The Use of Time Series Processes to Model the Error Structure of Earnings in Longitudinal Data Analysis," *Journal of Econometrics*, Vol. 18, No. 1, pp. 83–114.
- Meghir, Costas, and Luigi Pistaferri, 2004, "Income Variance Dynamics and Heterogeneity," *Econometrica*, Vol. 72, No. 1, pp. 1–32.
- Modigliani, Franco, and Shi Larry Cao, 2004, "The Chinese Saving Puzzle and the Life Cycle Hypothesis," *Journal of Economic Literature*, Vol. 42, pp. 145–70.
- Moffitt, Robert, and Peter Gottschalk, 1995, "Trends in the Transitory Variance of Male Earnings in the U.S, 1970-1987," Mimeo, Johns Hopkins University.
- Moffitt, Robert, and Peter Gottschalk, 2009, "Trends in the Transitory Variance of Male Earnings in the U.S., 1970-2004," 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*, Vol. 13, No. 3, pp. 305–14.
- Qian, Yingyi, 1998, "Urban and Rural Household Saving in China," *IMF Staff Papers*, Vol. 35, No. 4, pp. 592–627.
- Sin, Yvonne, 2005, "Pension Liabilities and Reform Options for Old Age Insurance," World Bank Working Paper No. 2005-1.
- Song, Michael, and Dennis Yang, 2010, "Life Cycle Earnings and the Household Saving Puzzle in a Fast-Growing Economy" Manuscript, Chinese University of Hong Kong.
- Topel, Robert H., and Michael P. Ward, 1992, "Job Mobility and the Careers of Young Men," *Quarterly Journal of Economics*, Vol. 107, No. 2, pp. 439–79.
- Violante, Giovanni, 2002, "Technological Acceleration, Skill Transferability, and The Rise in Residual Inequality," *Quarterly Journal of Economics*, Vol. 117, No. 1, pp. 297–338.
- Wei, Shang-Jin, and Xiaobo Zhang, 2009, "The Competitive Saving Motive: Evidence from Rising Sex Ratios and Savings Rates in China," NBER Working Paper No. 15093.

Table 1. Summary Statistics

Wave	Observations (Households)	Household	Labor Earnings (in RMB at 2006 prices)		Income (in RMB at 2006 prices)		At least some high school		Head employed by SOCE	
		Size Mean	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.
1989	493	3.8	7356.1	7656.7	12830.5	9171.0	0.36	0.48	0.81	0.39
1991	464	3.6	7403.7	3852.9	11537.1	5935.2	0.43	0.50	0.84	0.37
1993	336	3.5	9115.1	6199.9	13336.0	11002.0	0.48	0.50	0.84	0.37
1997	406	3.4	11279.0	7535.0	14980.6	9230.9	0.57	0.50	0.80	0.40
2000	444	3.2	16529.6	17334.7	21293.9	18831.4	0.67	0.47	0.77	0.42
2004	467	2.9	21934.5	18867.4	28247.5	25645.7	0.72	0.45	0.73	0.45
2006	477	2.9	25683.7	32553.3	32042.5	34845.5	0.72	0.45	0.64	0.48

Notes: Based on an unbalanced panel of urban households from the China Health and Nutrition Survey, with household heads who are aged 25-59, not a student, with complete age and education information, and not reporting positive income from farming and raising livestock.

Table 2. Estimated Variance of Permanent and Transitory Shocks to Urban Household Income

2A. Permanent Variance									
Year	Family Income						Family Income Excl. Trans. and Sub.	Family Income Excl. Trans. , Sub. and Other Inc.	Family Income Excl. 1989 Data
	All	SOCE	Young	Old	Less Edu.	More Edu.	All	All	All
1993	0.023	0.017	0.020	0.020	0.032	0.005	0.030	0.035	-
	(0.010)	(0.009)	(0.013)	(0.025)	(0.028)	(0.009)	(0.013)	(0.014)	-
1997	0.024	0.020	0.015	0.023	0.272	0.018	0.029	0.030	0.029
	(0.010)	(0.008)	(0.015)	(0.018)	(0.027)	(0.009)	(0.011)	(0.012)	(0.012)
2000	0.009	0.023	0.017	0.000	0.000	0.015	0.017	0.008	0.010
	(0.012)	(0.015)	(0.015)	(0.015)	(0.022)	(0.011)	(0.014)	(0.011)	(0.014)
2004	0.018	0.022	0.021	0.014	0.073	0.008	0.024	0.023	0.017
	(0.009)	(0.012)	(0.011)	(0.017)	(0.044)	(0.009)	(0.010)	(0.009)	(0.011)

2B. Transitory Variance									
Year	Family Income						Family Income Excl. Trans. and Sub.	Family Income Excl. Trans. , Sub. and Other Inc.	Family Income Excl. 1989 Data
	All	SOCE	Young	Old	Less Edu.	More Edu.	All	All	All
1991	0.040	0.029	0.053	0.027	0.025	0.045	0.059	0.046	-
	(0.012)	(0.011)	(0.019)	(0.013)	(0.016)	(0.013)	(0.019)	(0.018)	-
1993	0.085	0.061	0.086	0.104	0.117	0.103	0.115	0.103	0.075
	(0.017)	(0.017)	(0.024)	(0.038)	(0.072)	(0.026)	(0.021)	(0.022)	(0.017)
1997	0.090	0.066	0.095	0.129	0.141	0.061	0.084	0.105	0.081
	(0.036)	(0.022)	(0.045)	(0.078)	(0.105)	(0.028)	(0.038)	(0.033)	(0.038)
2000	0.153	0.131	0.168	0.119	0.164	0.152	0.182	0.182	0.165
	(0.026)	(0.025)	(0.032)	(0.055)	(0.111)	(0.031)	(0.027)	(0.029)	(0.029)
2004	0.162	0.149	0.152	0.193	0.111	0.169	0.181	0.129	0.160
	(0.030)	(0.028)	(0.0351)	(0.051)	(0.100)	(0.032)	(0.030)	(0.025)	(0.027)


Notes: Variance estimates based on the decomposition described in Section III. “All” refers to the full sample. SOCE subsample restricted to those where head was an SOCE employee when the household first appeared in the sample. Young and old households are defined as those whose household heads were below and above 40 years of age, respectively. “Less educated” is the group of households whose head does not have a high school degree. The last three columns correspond to (i) family income excluding subsidies and transfers; (ii) family income excluding subsidies, transfers and other income; and (iii) family income excluding data from year 1989. “Other income” includes income from leased land (only for 1989), rent from non-land assets, rent from lodgers or boarders, and other unspecified sources of income.

Table 3. Robustness Checks: Simulated Rise in Savings Using Various Preference Parameters

Robustness Checks: Simulated Rise in Savings Using Various Preference Parameters	Risk Aversion	Discount Factor	Jump in saving rates at time of shock for different ages of hhold. head						
			25	30	35	40	45	50	55
Baseline assumptions	4.5	0.97	5.4%	5.3%	5.4%	5.1%	5.8%	6.6%	7.8%
Larger drop in replacement rate post-reform	4.5	0.97	5.5%	5.6%	6.4%	7.2%	9.0%	10.6%	12.8%
Baseline drop in replacement rate with alternative preference parameters	3.5	0.97	5.0%	4.9%	5.1%	5.3%	6.8%	7.8%	8.9%
	2.0	0.97	4.4%	4.4%	4.4%	4.2%	5.0%	6.8%	9.5%
	4.5	0.95	5.3%	5.2%	5.3%	5.5%	6.3%	7.0%	8.2%
	3.5	0.95	4.9%	4.9%	5.0%	5.0%	6.4%	8.0%	9.3%
	2.0	0.95	4.2%	4.3%	4.3%	4.1%	4.6%	5.6%	8.5%
	4.5	0.98	5.5%	5.2%	5.2%	4.7%	5.6%	6.4%	7.6%
	3.5	0.98	5.1%	5.0%	5.2%	5.6%	6.9%	7.5%	8.7%
	2.0	0.98	4.5%	4.4%	4.4%	4.3%	5.3%	7.5%	9.9%
Preferences calibrated based on mean saving rate of 5-year age groups	7.5	0.99	2.2%	1.8%	2.2%	2.3%	3.3%	4.1%	5.1%
	8.1	0.97	2.1%	1.6%	2.1%	2.1%	3.1%	3.9%	5.0%


Notes: The first row reports the change in saving behavior when the transitory variance of income increases from 0.04 to 0.08 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 replacement rate post-reform is larger (declines to 0.50). 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.

Figure 1. Household Saving Rates




Source: National Bureau of Statistics, Flow of Funds data and Urban and Rural Household Survey. Saving rate from national accounts is significantly higher than that from the household surveys. This discrepancy is common (it is present in most countries), and can be due to differences in definitions of income and consumption, methodology and sample coverage.

Figure 2. Gross Saving Rates by Sector




Source: National Bureau of Statistics (Flow of Funds data)

Figure 3. Urban Household Saving Rates by Age of Household Head



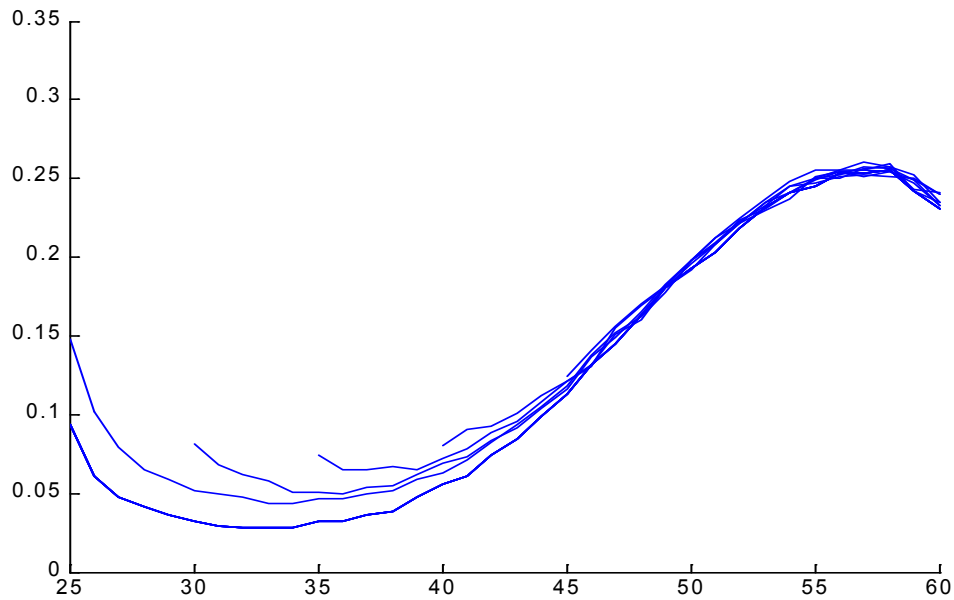
Notes: Based on a 10 province/municipality subsample of the National Bureau of Statistics Urban Household Surveys. Saving rates are smoothed by a moving average with 4 neighboring age averages. For details on the data and how saving rates are defined, please see Chamon and Prasad (2010).

Figure 4. Labor Market Turnover: Annualized Transition Probabilities




Note: 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.

Figure 5A. Estimated Age Profile of Saving Rates Before and After Rise in Variance of Transitory Income Shock




Notes: Dashed line corresponds to the saving behavior when the variance of transitory income shocks is 0.04. Other lines indicate the saving behavior when that variance is 0.08 if the change were to occur when the household was at that respective age.

Figure 5B. Jump in Saving Rate Following Rise in Variance of Transitory Income Shock




Notes: Each line corresponds to the jump in the saving rate in Figure 5A after the increase in the variance of transitory income shocks for a household with household head at that age.

Figure 6A. Estimated Age Profile of Saving Rates Before and After Pension Reform




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 household head was at that respective age.

Figure 6B. Jump in Saving Rates Following Pension Reform




Notes: Each line corresponds to the jump in the saving rate in Figure 6A after pension reform for a household with household head at that age.

Figure 7A. Estimated Age Profile of Saving Rates Before and After Rise in Transitory Variance of Income and Pension Reform




Notes: Dashed line corresponds to the saving behavior when the variance of transitory income is 0.04 and $\eta=0.75$. Other lines correspond to saving behavior when the variance of transitory income is 0.08 and $\eta=0.60$ if the change were to occur when the household head was at that respective age.

Figure 7B. Jump in Saving Rates Following Rise in Transitory Variance of Income and Pension Reform



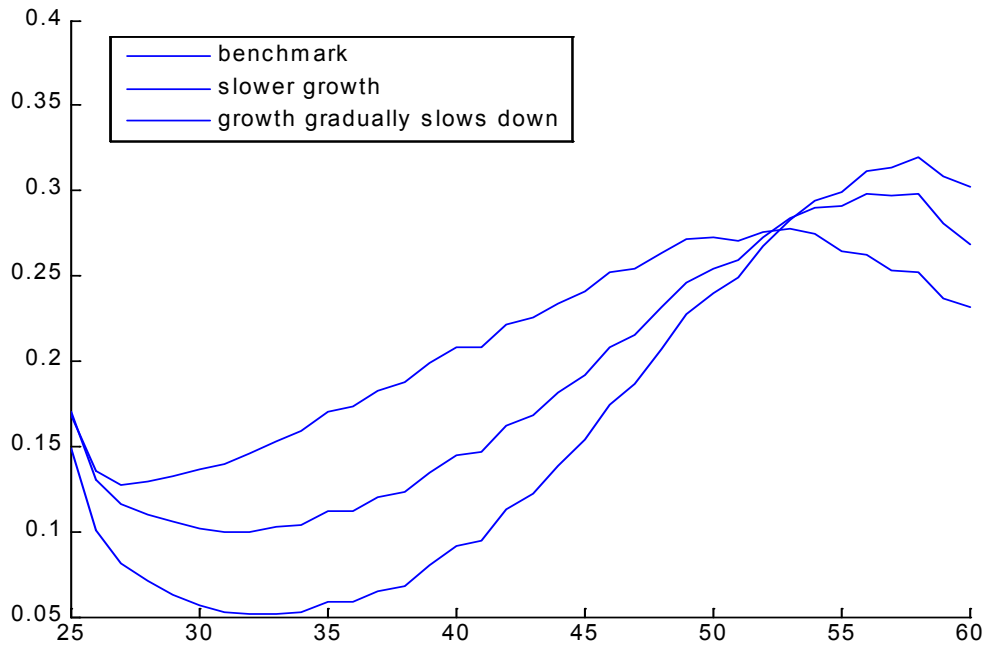
Notes: Each line corresponds to the jump in the saving rate in Figure 7A after the shock and pension reform occurs for a household with household head at that age. Changes in saving rates are smoothed by a moving average with 3 neighboring age averages.

Figure 8. Projected Cross-Sectional Change in Saving Rates After Rise in the Transitory Variance of Income and Pension Reform



Notes: Lines correspond to the cross-sectional age profile of savings implied by the profiles in Figure 7.

Figure 9. Estimated Age Profile of Saving Rates for a Household with 25 Year Old Head, Under Different Expected Income Growth Paths



Notes: Plot traces the evolution (at different ages) of expected saving rates for a household with a 25 year old head under our baseline parameter values (solid line), under a growth path where expected income growth is 1 percentage points lower every year than in our baseline (dotted line) and in a scenario where expected income growth is flattening at a rate of 0.1% every year (but never declines below the lowest value in the baseline path) (dashed line).