Aggregate Fluctuations in a Quantitative Overlapping Generations Economy with Unemployment Risk

Aubhik Khan

Ohio State University

February 2016

ABSTRACT

We explore business cycles in quantitative overlapping generations economies where households face both unemployment risk and, conditional on employment, income risk. These life-cycle economies exhibit large inter-generational differences in average wealth which, given estimated processes for income risk, allow them to reproduce much of the wealth inequality in the data.

In this environment, we examine the effect of aggregate shocks on the distribution of consumption, income and wealth across households and how this distribution shapes the aggregate response to these shocks. While TFP fell relatively little, in the recent US recession, there was a large fall in hours worked. Consistency with these observations, in the model, requires a substantial and persistent increase in unemployment risk that disproportionately affects younger working households. Conversely, explaining the reduction in aggregate consumption implies a financial shock to households' net return on assets. This drives a non-monotonic welfare cost of the recession, over generations, varying by the share of income derived from capital and labour. Overall, young workers, and older workers near retirement, suffer the most from the increase in unemployment and the fall in the return on wealth, respectively.

Inequality in wealth also shapes the aggregate response of the economy. Aggregate investment falls by more when households face little income risk, holds less precautionary savings, and are more responsive to changes in real interest rates.

mail@aubhik-khan.net.

1 Introduction

The recent recession involved a large fall in GDP and a persistent reduction in employment alongside a comparatively small decrease in measured TFP. We examine these movements through a quantitative overlapping generations economy where households face both employment and income risk, and there is both intra- and inter-generational inequality in income and wealth. In this environment, the small change in TFP and the large decrease in hours worked explains the observed fall in GDP and the most of the of decline in business fixed investment.

In contrast to a typical incomplete markets model, generations vary in their savings rates and in their shares of income from labour and capital. As consistency with the last US recession requires a sharp rise in unemployment risk, we find a disproportionate impact on younger, less wealthy households. In contrast, older households are more insulated against a fall in labour income as they tend to have higher levels of wealth, allowing them to consumption, or have retired.

Shocks to TFP and employment are insufficient to reproduce the fall in aggregate consumption seen over the 2007 recession. However, a financial shock that reduces households' return on assets allows the model to reproduce the large drop in detrended consumption. This drives welfare costs that are non-monotone in age; the reduction in consumption is greatest for the households that have recently joined the labour force, then fall with labour market experience only to rise again for workers nearing retirement. The young suffer the most because of their low levels of wealth and the importance of labour earnings in their income. Older workers have accumulated, on average, substantial wealth and thus have a higher share of income from capital. The reduction in their capital income reduces their consumption relative to households that have accumulated less wealth. While employment and financial shocks are essential in delivering a recession, in the model economy, that captures the data, the shock to TFP is inessential for understanding this recent economic episode.

Inequality of consumption, income and wealth, by age and within age groups is not only affected by shocks, but also shapes the aggregate response to these shocks. As an example, realistic levels of income risk drive a strong precautionary savings motive which makes households relatively unresponsive to changes in real interest rates. When we reduce income risk, households decrease savings and aggregate investment experiences a larger drop over the recession. The model somewhat overpredicts movements in investment and realistic levels of income risk are important in reducing the distance between model and data.

In our model economy, households face shocks to income that have both a highly persistent and a transitory component, both of which are estimated from the PSID. As in a standard incomplete markets model with infinitely-lived households, they may borrow or save and there is a borrowing limit. Young households are wealth-poor and, on average, have high savings rates. Wealth rises over the life-cycle until retirement. This retirement savings motive generates substantial additional inequality, over and beyond what may be explained by income risk and the calibrated model generates a Gini coefficient for wealth over 70 percent. As they have accumulated little savings, younger households have a higher share of income from labour, and thus larger lifetime effects of persistent unemployment on consumption and subsequent wealth. Older households have, on average, high wealth to income ratios and are better able to absorb unemployment risk if they are still working. However this makes them more sensitive to financial shocks that decrease their return on assets. Moreover, whether to labour or capital income, have persistent effects on consumption as households have finite lives.

A new literature explores the role of changes in the distribution of households or firms in the amplification and propagation of aggregate series over the business cycle.[1] Krueger, Mitman and Perri (2015) show that realistic income and wealth inequality can amplify the fall in consumption, and thus GDP, following a large TFP shock. Thus, while the seminal work of Krusell and Smith (1999) found that poor households were relatively unimportant in determining the evolution of aggregate capital, Krueger et al. find that such households are important in determining movements in aggregate consumption spending and thus GDP following large shocks to aggregate TFP. An important difference between Krueger et al. and this work is our overlapping generations setting. We explore the importance of heterogeneity in an environment where households differ not only in wealth, but also in age. Aggregate shocks, which imply equilibrium movements in real wages and interest rates, have different effects across generations given differences in their shares of income from capital and labour. Second, we assume a small shock to TFP, consistent with the 2007 recession, and introduce a shock to unemployment risk. The latter has strong distributional effects that drive sharp differences in the impact of the recession on different generations.

Glover et al. (2014) study the redistributional effect of shocks to asset prices in an overlapping generations model. In contrast, we explore the relation between inequality and movements in aggregate GDP and other series during a recession. Thus we endogenise

[1] Studies exploring the role of heterogeneous households include Krueger, Mittman and Perri 2014, Guerrieri and Lorenzoni 2015 and Glover, Heathcote, Krueger and Rios-Rull 2014. The aggregate effect of changes in the distribution of firms or entrepreneurs, following credit shocks, is studied in Khan and Thomas 2013 and Buera and Moll 2014.

production and, while we do not model asset prices, we allow for capital accumulation. In our setting, as younger households are relatively poor, their share of aggregate consumption, while more sensitive to earnings fluctuations driven by increases in unemployment, is small. As a result, increases in unemployment do not explain the large observed fall in aggregate consumption. As mentioned above, this difference between model and data is resolved using a persistent shock to households' net return on capital income. The addition of this financial shock implies large welfare costs that are non-monotonic across households of different ages. Thus our focus is on studying the distributional effects of aggregate shocks that explain movements in real series instead of asset prices; we find that the wealth of young and old workers falls the most over the recession, while middle-aged workers and retirees suffer smaller declines. In related work, Guerrieri and Lorenzoni (2015) explore the effect of shocks to households' borrowing limits on the persistence and size of a recession. They find that nominal rigidities and the zero lower bound for nominal interest rates are important in amplifying the real effects of shocks to household borrowing limits.

Most of our results are derived by solving perfect foresight paths in general equilibrium. However we also develop a solution method for the model with aggregate uncertainty. An independent contribution of this work is to show how to solve for stochastic equilibria in overlapping generations models with uninsurable intra-generational risk. Krueger and Kubler (2010) find that aggregate state space approximation, as developed by Krusell and Smith (1999) where the mean of the distribution of capital is used as a proxy for the aggregate state, may imply large Euler equation errors when applied to overlapping generations models. In our model, large employment risk shocks imply important changes in the distribution of wealth within and across generations. This makes the use of aggregate capital, and a parametric forecasting rule, problematic. Additionally, the presence of uninsurable earnings risk implies that an exact solution using Smolyak polynomials, as in Krueger and Kubler, is infeasible. We solve for stochastic equilibria by extending Reiter's (2010) backwards induction method with an approximate aggregate state that is a vector of moments from the distribution of wealth, each element of which is the total capital held by several consecutive generations. Our implementation relies on using a sparse, correlated grid of aggregate moments for this potentially large aggregate state vector. After solving the model, and updating the distribution at each grid point using results from a simulation, we also update the grid itself. This provides a solution method for overlapping generations models with many generations and uninsurable income risk.

Section 2 of the paper describes the model, section while section 3 describes the calibration of the model economy including the estimation of income shocks. Section 4 contains results and section 5 concludes.

2 Model

The economy has a unit measure of households in overlapping generations. Each generation has a life-cycle characterised by an initial period of work followed by retirement. In the working part of their lives, households face unemployment risk as well as income risk when employed. Unemployed households receive benefits and there are public pensions for retired households. Households may accumulate wealth or borrow, subject to a limit. There is a representative firm that employs all working households, paying each a wage proportional to their efficiency units of labour, and rents capital from households.

2.1 Households

Households that have been working j years have experience $l\left(j\right)$. Each household has a persistent idiosyncratic labour productivity $\zeta_k \in \left\{\zeta_1, \ldots, \zeta_{N(\zeta)}\right\}$ that has transition matrix, $\pi_{ij} \geq 0$, $i, j = 1, \ldots, N\left(\zeta\right)$, where $\sum_{j=1}^{N(\zeta)} \pi_{i,j} = 1$ for $i = 1, \ldots, N\left(\zeta\right)$.

Households faces a risk of unemployment which is uncorrelated with their labour productivity. Let $e \in \{0, 1\}$ describe employment when $e = 1$ and unemployment otherwise. The transition probabilities between employment and unemployment are $\pi_{ij}^{e}\left(s\right) \geq 0$, $i = 1, 2$ and $j = 1, 2$, with $\sum_{j=1}^{2} \pi_{ij}^{e}\left(s\right) = 1$. The exogenous component of the aggregate state is s, and unemployment risk in the economy may change with aggregate shocks as in Krusell and Smith (1999). Both labour productivity shocks, and unemployment shocks, are independent of a household's working age, j. Households retire after $J_r - 1$ years.

The individual's state is (j, a, e, ζ) where $a \in [\underline{a}, \infty)$ are assets and $j \in \mathbf{J} = \{1, \ldots, J\}$ describes the years since the start of work. The budget constraint for a working household, $j = 1, \ldots, J_r - 1$, is

$$\begin{aligned} c + a' &\leq (1+r)\,a + (1-\tau)\,w\zeta l\left(j\right) e + (1-e)\,b_0\zeta, \\ c &\geq 0,\ 0 \leq n \leq 1 \text{ and } a' \geq \underline{a}. \end{aligned}$$

Above, c is current consumption, r is the real interest on savings and w is the real wage, both of which are functions of the aggregate state. When employed, households receive $(1-\tau)\,w\zeta l\left(j\right) e$ where τ is the labour income tax rate. Unemployed households receive benefits, $b_0\zeta$ that are proportional to their labour productivity. The persistence of labour productivity implies that unemployment benefits vary with actual earnings over previous employment. While households may borrow, their savings cannot cannot fall below the borrowing limit, $\underline{a} \leq 0$.

Labour income taxes τ finance government spending, unemployment benefits and social security payments to retired households, b. These households, in cohorts $j = J_r, \ldots, J$, have the following budget constraint,

$$\begin{aligned} c + a' &\leq (1+r)\,a + b\zeta, \\ c &\geq 0 \text{ and } a' \geq \underline{a}. \end{aligned}$$

Above, ζ represents the household's labour productivity in its last period of work. Assuming social security benefits are proportional to labour income in the last year of work, we allow them to vary with lifetime earnings.

A working household of generation $j = 1, \ldots, J_r - 1$ solves

$$\begin{aligned} v\left(j, a, e_h, \zeta_k, s_m, A\right) &= \max_{c,a'} \Big(u\left(c\right) \\ &\quad + \beta \sum_{m=1}^{N(s)} \pi^s_{mn} \sum_{i=1}^{2} \pi^e_{hi}\left(s_n\right) \sum_{l=1}^{N(\zeta)} \pi_{kl} v\left(j+1, a', e_i, \zeta_l, s_n, A'\right) \Big) \\ &\quad \text{subject to} \\ c + a' &\leq (1+r)\,a + (1-\tau)\,w\zeta l\left(j\right) e + (1-e)\,b_0\zeta \\ c &\geq 0,\ a' \geq \underline{a}. \end{aligned} \tag{1}$$

The aggregate state evolves over time according to $A' = \Gamma\left(s, A\right)$, and the real interest rate, r, the real wage, w, as well as unemployment benefits, b_0, are functions of the aggregate state (s, A). The exogenous component of the aggregate state, s, is a Markov Chain, $\{\pi^s_{mn}\}_{m,n=1}^{N(s)}$ with $\pi^s_{mn} \geq 0$ and $\sum_{n=1}^{N(s)} \pi_{mn} = 1$ for each $m = 1, \ldots, N\left(s\right)$

We describe retired households starting with the last period, when households simply consume their cash on hand.

$$v\left(J, a, e, \zeta_k, s_m, A\right) = u\left(\left(1+r\right)a + b\zeta_k\right). \tag{2}$$

Social security payments, per unit of labour productivity, b, are also a function of the aggregate state. In earlier periods, retired households in cohorts $j = J_r, \ldots, J-1$ solve the following borrowing and savings problem.

$$
\begin{aligned}
v\left(j,a,e,\zeta_k,s_m,A\right) &= \max_{c,a',n}\Big(u\left(c\right)+\beta\sum_{m=1}^{N(s)}\pi^s_{mn}v\left(j+1,a',e,\zeta_k,s_m,A\right)\Big) \qquad (3)\\
&\quad \text{subject to}\\
c+a' &\leq \left(1+r\right)a+b\zeta_k\\
c &\geq 0,\ a'\geq \underline{a}.
\end{aligned}
$$

Retired households' income will not vary with e, and ζ_k does not change over time. We include both terms in their value functions for consistency with (1) when $j=J_r-1$. However, the only source of uncertainty for retired households is aggregate risk.

2.2 Production and government

A representative firm uses capital and labour to produce output. We begin by describing the aggregate supply of efficiency units of labour. Let π^e_1 denote the fraction of employed households which, in stationary equilibrium, would be the invariant distribution implied by $\{\pi^e_{hi}\left(\overline{s}\right)\}^2_{h,i=1}$ where $\overline{s}$ is the steady state value for the exogenous aggregate state. Further, $\{\pi^0_k\}^{N(\zeta)}_{k=1}$ describes the invariant distribution of households over labour productivity, determined by $\{\pi_{kl}\}^{N(\zeta)}_{k,l=1}$. Lastly, let $\left\{\mu^a_j\right\}^J_{j=1}$ describe the number of households in each cohort $j=1,\ldots,J$. The aggregate labour supply is the sum of working age households, over labour productivity, that are employed,

$$L=\pi^e_1\sum_{j=1}^{J_r-1}\mu^a_j\sum_{k=1}^{N(\zeta)}\pi^0_k\zeta_k.$$

In the following, we assume that $\mu^a_j=\frac{1}{J}$ for all j. Let $\mu\left(j,e,\zeta,a\right)$ describe the measure of households of working age j with labour productivity ζ, employment status, e and wealth a. The aggregate capital stock is

$$K=\sum_{j=1}^{J_r}\int_{\left\{\zeta_1,\ldots,\zeta_{N(\zeta)}\right\}\times[\underline{a},\infty)}a\mu\left(j,d\left[e\times\zeta\times a\right]\right).$$

Given capital and labour, production is

$$Y=zK^{\alpha}L^{1-\alpha}$$

where z is aggregate total factor productivity and $0<\alpha<1$.

In equilibrium, the real return on savings, and the real wage, are given by

$$r+\kappa = \alpha z K^{\alpha-1}L^{1-\alpha}-\delta \tag{4}$$
$$w = (1-\alpha) z K^{\alpha}L^{-\alpha}, \tag{5}$$

where $\delta \in (0,1)$ is the rate of depreciation of capital. A financial shock, $\kappa \geq 0$, may reduce the real return on savings, r, below that marginal product of capital. We assume that κ is 0 in the steady state. The exogenous aggregate state is $s=(z,u,\kappa)$ where u is the component that determines employment risk for households.

As already mentioned, tax revenues are used to fund social security, unemployment benefits and government spending. The government budget constraint is

$$b\sum_{k=1}^{N(\zeta)}\pi_k^0\zeta_k\sum_{j=J_r}^{J}\mu_j^a+(1-\pi_1^e)\,b_0\sum_{k=1}^{N(\zeta)}\pi_k^0\zeta_k\sum_{j=1}^{J_r-1}\mu_j^a+G=\tau wL \tag{6}$$

where G is government spending.

2.3 The distribution of households

We assume that initial wealth of newly working households is 0, and that the distribution of employment and labour productivity, when $j=1$, is given by π_1^e and π_k^0. Given $\mu(1,e,\zeta,a)$, the distribution of newly working households over employment, labour productivity and wealth, the evolution of households over age is given by

$$\mu(j+1,e',\zeta_k,B)=\sum_{e=0}^{1}\pi_{ee'}\sum_{i=1}^{N(\zeta)}\pi_{ik}\int_{\{(e,\zeta_i,a)|g(j,e,\zeta_i,a)\in B\}}a\mu(j,e,\zeta_i,a)$$

where $g(j,e,\zeta,a)$ is the savings decision rule for a household of age j with current employment status $e\in\{0,1\}$, labour productivity ζ and beginning of period wealth, a.[2] The distribution $\mu(j,e,\zeta,a)$ is the endogenous component of the aggregate state, A.

3 Calibration

We set the length of a period to one year, and assume that households begin working at age 25. Given this, we set $J=60$ and assume that retirement is at age 65, implying $J_r=40$. We assume that period utility takes the form $u(c)=\frac{c^{1-\sigma}}{1-\sigma}$ and set the coefficient of relative risk aversion, $\sigma=2$. The subjective discount factor, $\beta=0.955$ implying a 4

[2] This decision rule solves (1) when $j=1,\ldots,J_r-1$ and (3) for $j=J_r,\ldots,J-1$.

percent real interest rate. The remaining parameters take common values. The annual depreciation rate is $\delta = 0.069$ and capital's share of output is $\alpha = 0.36$. The average labour tax rate is $\tau = 0.3$. Unemployment benefits are 30 percent of annual earnings, $b_0 = 0.3w$. Given the annual calibration of the model, this is chosen to represent, on average, six months of benefits at 60 percent of labour income, $b_0\zeta$, when employed. The social security replacement rate is $b = 0.45$.

3.1 Earnings Shocks Estimation

The estimated earnings shock process here is taken from Khan and Kim (2014) which itself is closely related to the approach in Heathcote, Storesletten and Violante (2010). First, labor market experience and the earnings shock process is estimated using PSID data between 1968 and 2011. The PSID data is a longitudinal survey of a sample of US individuals and families conducted annually from 1968 to 1997, and biennially since 1997. The original 1968 PSID sample combines the Survey Research Center (SRC) and the Survey of Economic Opportunities (SEO) samples. We use the U.S. population representative SRC sample.

Household earnings are defined as the sum of the annual earnings of the head of household and spouse. Total annual earnings include all income from wages, salaries, bonuses, overtime, commissions, professional practice and the labor part of farm and business income. We select households with no missing values for education and self-employment status where: 1) the head of households' age is between 25 and 59 years old, 2) neither the head or spouse has positive labor income but zero annual hours worked, 3) the hourly wage is not less than half of the minimum wage, 4) income is not from self-employment. Earnings are deflated with the CPI and expressed in 2013 dollars. After the sample selection, there are only 23 top-coded observations out of 42,127 observations. Following Autor and Katz (1999), we multiply all top-coded observations by a factor of 1.5 times the top-coded thresholds.

Let $y_{i,j,t}$ be the earnings of household i with head at age j in year t. We run an OLS regression of log earnings on time dummies; an interaction term with education, edu_i, and time dummies, labor market experience, θ; and experience-squared, θ^2.[3]

$$\log y_{i,j,t} = \beta_{t,0} + \beta_{t,1} edu_i + \beta_3 \theta + \beta_4 \theta^2 + \widehat{r}_{i,j,t}$$

[3] Labor market experience is measured as age minus years of schooling minus 5. In years missing the variable for years of schooling, we proxy years of schooling for the individuals with a college degree as 16 and for the individuals without a college degree as 12.

Efficiency units of labor depend on both labor market experience, $l(j) = e^{(\beta_3\theta+\beta_4\theta^2)}$, and idiosyncratic productivity shocks. We estimate $\beta_3 = 0.04099$ and $\beta_4 = -0.00067$, this leads to an evolution of experience that increases earnings through the first 27 years of working life, when average earnings rise 71 percent relative to their initial level. Thereafter, total earnings fall until retirement.

The regression residuals $\widehat{r}_{i,j,t}$ are assumed to be the sum of idiosyncratic productivity shocks, $\log \epsilon_{i,j,t}$, and measurement error, $\widetilde{v}_{i,j,t}$.[4] Idiosyncratic productivity shocks consist of both a persistent component, $\log \epsilon^p$, and transitory component, $\log \epsilon^v$. To be specific,

$$\log \epsilon_{i,j,t} = \log \epsilon^p_{i,j,t} + \log \epsilon^v_{i,j,t}$$
$$\log \epsilon^p_{i,j,t} = \rho \log \epsilon^p_{i,j-1,t-1} + \eta_{i,j,t}$$

where $\eta_{i,j,t} \sim N(0, \sigma^2_{\eta_t})$ and $\log \epsilon^v_{i,j,t} \sim N(0, \sigma^2_{vt})$.

Following Heathcote, Storesletten, and Violante (2010), we estimate year-varying shock variances $\{\sigma^2_{\eta_t}, \sigma^2_{vt}\}$, the persistence of the shock $\{\rho\}$, and the variance of persistent shock for initial age σ^2_π using minimum distance methods. We use survey data from 1968 to 2011, but only estimate the variances up to 2009 because of the finite sample bias at the end of sample period. As Heathcote et al.. (2010) point out, although the variance of the persistent shock for the missing years can be theoretically pinned down by the available information from adjacent years, the resulting estimates for missing years are downward-biased because of insufficient information. Therefore, we follow their approach of Heathcote et al. (2010) for the estimates in the missing years by taking the weighted average of the two adjacent years.

The estimated persistence of earning shock, ρ, is 0.9856. The benchmark model economy is calibrated to 2004 with a smoothed estimate of the persistent shock variance, 0.0363, and a transitory shock variance of 0.1913. We abstract from the estimated variance of the initial distribution of the persistent shock, which is estimated to be 0.1450, and instead allocate households, in their first period of work, using the standard deviation implied by the invariant distribution of the persistent shock.

Unemployment risk is a discrete stochastic process with the probability of remaining unemployed set to 0.1, and the probability of losing employment set to 0.05. The annual calibration of the model makes it difficult to model unemployment duration, and the expected duration of unemployment, 1.1 years, is about twice that seen in the data. However, the overall unemployment rate in the stationary state, is a little over 5 percent.

[4]Given the similarity in the estimates of earnings process to wage process estimated in Kim (2015), we use French (2002)'s estimate of a variance of measurement error in log hourly wages of 0.02.

3.2 Solution

Individual decision rules are solved using the endogenous grid method given the aggregate state (Carroll 2005). The Rouwenhorst algorithm is used to combine the transitory shock, using 2 grid points, and the persistent shock, using 5 grid points. We solve household value function starting with $j = J$, on a fixed grid of assets, productivity and employment state using (**??**). When using the perfect foresight solution method, we solve for value functions at each date t. Given $v(J, a, e, \zeta_k, t)$, we determine the value function, for earlier ages, $j = 1, \ldots, J-1$ working backwards using (3) for j greater than J_r, the retirement age, and (1) for j less than J_r. At each age, j, we assume $a' = a_i \in \mathbf{A} = \{a_1, \ldots, a_N\}$, a fixed grid of values for assets (we set $a_1 = \underline{a}$) and the numerical derivative of $\sum_{i=1}^{2} \pi_{hi}^{e}(s_t) \sum_{l=1}^{N(\zeta)} \pi_{kl} v(j+1, a', e_i, \zeta_l, t+1)$ to infer current consumption and thus cash on hand. While cash on hand differs for retired or working age households, retirement is a function of j and we can compute current wealth, a, as a function of a', labour productivity, ζ, and employment, e. If a $j \geq J_r$, decision rules are independent of e. Working age households, even if unemployed will have savings that varies with ζ as the current income shock determines the distribution of future labour productivity if they obtain employment next period. After retirement the ζ in the last working period determines social security benefits. This determines an endogenous grid, $a(j, a', e, \zeta, t)$ and households' current expected discounted lifetime value upon this grid. We interpolate the value function and savings, $a' = a_i$, onto the original grid, $\mathbf{A}$ using a. Monotonicity of decision rules in current wealth implies that if a, given (j, a', ζ, e, t) is less than $\underline{a}$ then we can ignore it. Instead if $a(j, a', \zeta, e, t) > \underline{a}$ when $a' = \underline{a}$, then all levels of current assets a with $\underline{a} \leq a < a(j, \underline{a}, \zeta, e, t)$ also have savings of $\underline{a}$. When implementing this method, we set $N = 200$, $a_{N\ =150}$ and $a_1 = -0.3$. The latter value is chosen to reproduce the share of households with negative assets.

Given decision rules, we iterate on the aggregate capital stock over a finite simulation, assuming the economy returns to the steady state in $T = 120$ periods. Any vector of capital, $\{k_t\}_{t=1}^{T-1}$ implies wages and interest rates that allow for the solution of households' value function and decision rules in a backwards loop. Next, these decision rules, alongside an initial distribution of households, is used to compute a supply of aggregate capital by households over each date. We iterate until the assumed vector of capital, which generates the real interest and wage each period, is consistent with the capital stock supplied by households. The initial distribution of households is the same as that in the steady state, with the exception that 5 percent of working age households are randomly reassigned from employment to unemployment. We record the distribution of households on a weighted grid with 1000 values and the same endpoints as the grid used for decision rules. The distribution has dimension $60 \times 10 \times 2 \times 1000$. There is one such distribution for each period $t = 1, ..., T$.

We also solve the model with aggregate uncertainty using backwards induction as described in Reiter (2010). We modify this method to use a vector of moments for the approximate aggregate state, each element representing the total wealth of households across several generations. We find that using 3 moments, with the capital held by young and middle-aged workers, and retired households, is efficient. Our application of backwards induction requires explicit parameterisation of the stochastic process not only for TFP, but also for employment shocks and capital income shocks. We set employment shocks to be perfectly correlated with TFP shocks. However we assume that shocks to households' return on wealth are rare, and, in an initial example, use the stochastic process for credit shocks described by Khan and Thomas (2013). Household value functions are now functions of 2 exogenous aggregate state, 3 moments for the endogenous aggregate state, as well as individual labour productivity and employment shocks and savings.

Movements in the aggregate state are correlated, as an example, when the wealth of retirees is below its mean, the wealth of young households is not above its mean. Thus, in a departure from Reiter's method, we use a single 3 dimensional grid of correlated states for the aggregate state. This allows for efficient computation of recursive equilibria. After solving the model, we update the set of grid points, and the moment consistent distribution of households at each such grid point, by simulation (see Reiter 2010). The actual moments of the distribution of households in the simulation are used in an outermost loop to update the grid of value for the endogenous aggregate state. The results from the stochastic simulation are not yet available.

4 Results

4.1 Steady State

We begin by studying the stationary state of the economy. Households are born with 0 assets and an initial labour productivity drawn from the invariant distribution of the bivariate Markov Chain of persistent and transitory shocks. Over time they borrow or save, and their earnings grow with experience. Over their life-cycles, households accumulate wealth, then after they retire, they dissave. Thus younger households tend to be poorer than older working households, and the middle-aged are wealthier than retirees. This drives inequality in wealth that is absent in an infinite-horizon model where the sole source

Figure 1: The Lorenz Curve

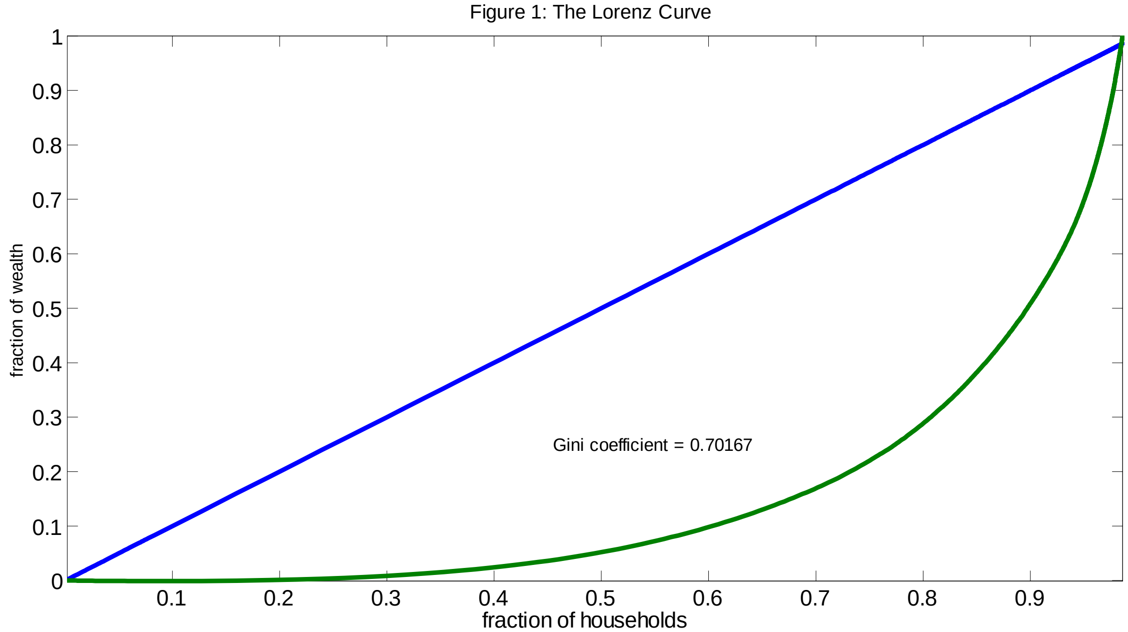

of difference in wealth across households is the result of differences in histories of labour productivity shocks.

Table 1: percentiles of the wealth distribution,
Gini coefficient and negative assets

	1	5	10	50	90	Gini	<0
SCF	0.29	0.51	0.64	0.97	1.0	0.77	0.09
model	0.11	0.37	0.53	0.95	1.0	0.70	0.08

Table 1 shows that the model exhibits far greater wealth inequality than would arise in an infinite-horizon model with the same distribution of income risk. The Gini coefficient for wealth is within 7 percentage points of the SCF data in 2004. While the model is calibrated to reproduce the share of negative assets, the absence of motives to accumulate wealth over and beyond consumption smoothing and retirement implies a skewness in the distribution of wealth that is less than seen in the data. Thus the wealthiest 1 percent of households hold 11 percent of total wealth in the economy, while, in the data, the corresponding number is 29 percent. A potential remedy would be greater skewness in earnings as in Diaz-Gimenez et al (2004) or, more generally, a distribution of labour productivity shocks that is substantially different from normal as reported by Guvenen et al. (2015). However, as seen in figure 1 the overlapping generations economy, for an empirically consistent process for income shocks, is able to generate inequality of wealth that is close to the data.

Figure 2 The Distribution of Wealth

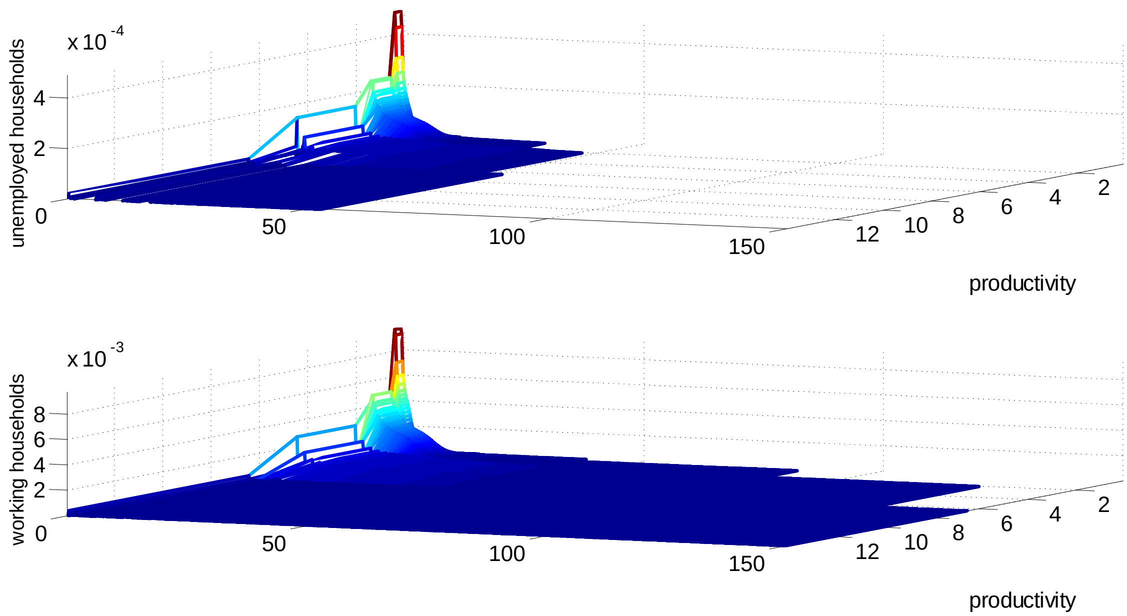

Figure 2 shows the distribution of wealth, as a function of income shocks, for unemployed and employed households, over all ages. The right horizontal axis represents labour productivity, while the front of each panel shows wealth. Clearly employment risk is negatively correlated with wealth. Moreover, there is considerable dispersion in the distribution of wealth, and as clear in the bottom panel, the persistence of income shocks implies that higher income households tend to be wealthier.

The accumulation of wealth over working lives, towards smoothing consumption after retirement, drives considerable inequality even in the absence of income risk. If we reduce the standard deviation of labour productivity shocks by 10, inter-generational wealth inequality remains and the Gini coefficient is still 0.49.

Figure 3 shows the average level of income, consumption and wealth in one generation as its members age. They retire after 45 years of work, and the estimated quadratic returns to labour market experience is responsible for the nonlinear trend in labour earnings seen in the top panel. As mentioned in the beginning of section 3, the social security replacement rate used in the model is 0.45, and retirement earnings are proportional to labour income in the last period of work. This income, averaged across households, is shown in the top panel starting with $j = 45$.

The equilibrium real interest rate in the steady state is 4 percent, while household's subjective discount factor is $\beta = 0.955$. Income risk and life-cycle savings drive wealth accumulation, as seen in the bottom panel. On average, households save until retirement,

Figure 3 A Cohort of 25000 Households

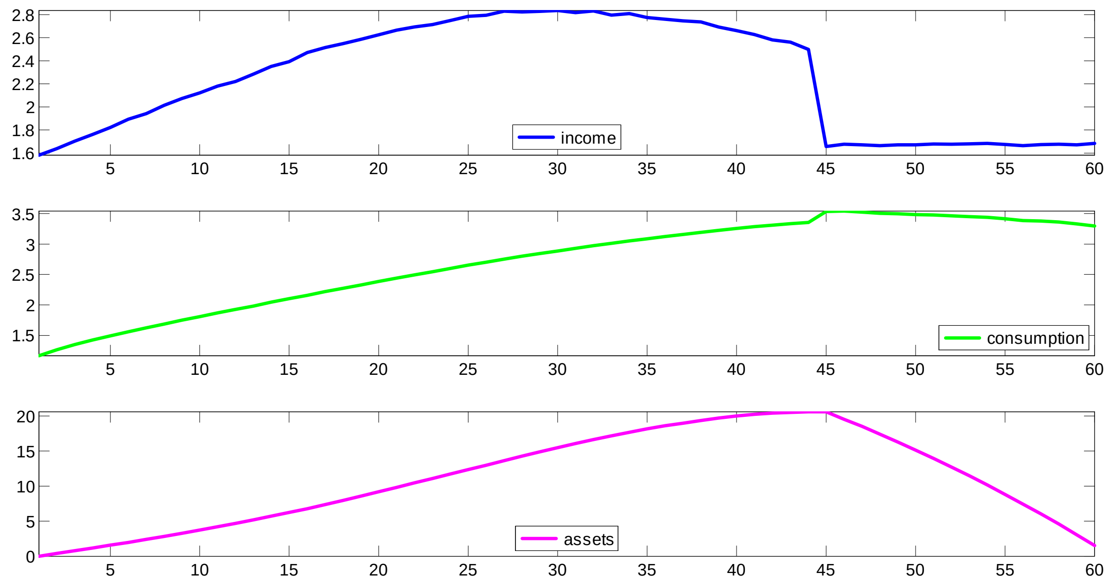

then gradually deccummulate their assets to smooth consumption relative to income. The rise in income and wealth implies a gradual increase in consumption while households are of working age. Consumption slowly falls over retirement, as households, facing no further risk, have a marginal rate of substitution less than one. The small rise in consumption at retirement is likely the result of a wealth effect, upon retirement, as households no longer need to maintain savings against idiosyncratic income risk.

Figure 4 reports intragenerational inequality, in the sample of households, using the cross-sectional standard deviations of income, consumption and wealth. As experience grows, income shocks imply greater earnings inequality. This starts to fall as the effect of the quadratic experience function on labour earnings start to reduce working age income, around $j = 30$ when households are 55 years old. Inequality in income results drives inequality in wealth as higher income households save more to both offset the effect of future shocks to income that may reduce consumption and towards retirement. As a result, we see the standard deviation of assets, across households, rise with working age in the bottom panel. As households near retirement, their savings rates fall, and thus both the growth in wealth, seen in the bottom of figure 3, and the rise in the standard deviation of wealth across households, in the bottom panel of figure 4, gradually slows as the cohort reaches $j = 45$.

The inequality generated by the model, solely from income risk propagated by life-cycle savings, suggests that it is a useful starting point to evaluate the role of income and wealth

Figure 4 A Cohort of 25000 Households with cross-sectional std.

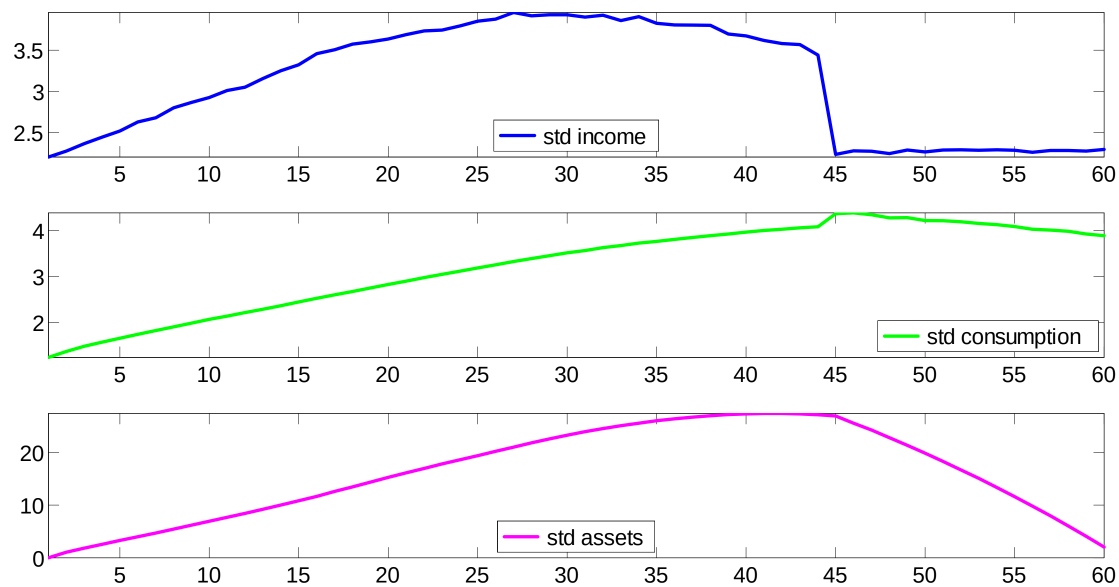

heterogeneity, and age, in determining the equilibrium response to aggregate shocks.

4.2 Business Cycles

The recession that began in 2007Q4 saw detrended GDP fall, through the trough in 2009Q2, by 5.69 percent. Consumption fell 4.08 percent, business fixed investment dropped by 18.98 percent and total hours worked fell 6.08 percent.[5] Surprisingly, measured TFP fell by only 2.18 percent over the same period. We explore the impact, on the model economy, of three shocks that work, to increasing effect in reproducing this Great Recession. We explore shocks to TFP, employment and capital income, across three cases.

The TFP shock has a persistence of 0.9, consistent with the Solow Residual measured by Khan and Thomas (2013), and is of the same magnitude as observed in the recession. We assume the same persistence for all shocks. The employment shock raises the unemployment rate to 10 percent from its ordinary level just above 5 percent. Further, it's propagated by assuming that households' transition probabilities for unemployment and employment change. Specifically, the probability of remaining unemployed rises by 5 percentage points, and the risk of losing employment increases by the same percentage. These shocks to the probabilities of remaining unemployed and losing employment have the same persistence as TFP, gradually returning to their steady state values.

Our first case studies the aggregate effect of TFP and employment shocks described

[5] See Khan and Thomas (2013) for details.

Figure 5 TFP, Employment and Capital Shocks

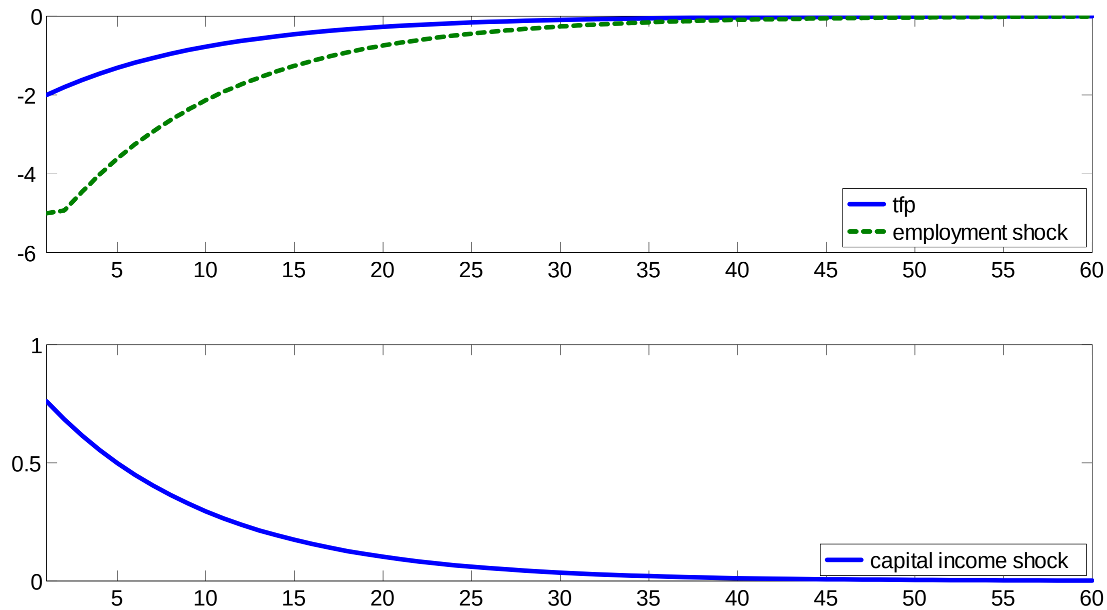

above. Our second case adds a shock to the real return on capital income, essentially households see a decline in the real interest they earn on assets, and this decline is chosen to reproduce the observed reduction in consumption over the recession, given the shocks to TFP and employment. This shock is implemented by adjusting the real interest rate received by households to $r_t - \kappa_t$ where firms rent capital at the real interest rate r_t. We assume that $\kappa_1 = 0.0076$ and that κ has the same persistence as TFP and the shock to unemployment. The third case eliminates all but the TFP shock, for comparison to the first and second case. Figure 5 plots in percent all three shocks that are used in various combinations across the three cases described below.

The baseline case, with TFP and employment shocks, is able to explain the decline in GDP and much of the fall in investment observed in the 2007 recession. GDP fall 6.9 percent upon impact, and investment drops 13.6 percent, as shown in Figure 6. However the fall in consumption is relatively small. In Figure 6a, we see it unravel in a familiar hump-shaped manner, eventually decreasing 2.8 percent 7 periods after the shock. This is about 70 percent of the decrease seen in the data.

The welfare effects of the recession, driven by shocks to total factor productivity and employment risk, vary across households. This is evident in Figure 6b which shows average consumption in each generation over the first 5 periods from the start of the recession. These cross-sectional consumption profiles are percent deviations from their steady state values. While lifetime welfare is not equivalent to five years consumption, it provides an

Figure 6a GDP, Cons., Inv. and Capital with TFP and Employment Shocks

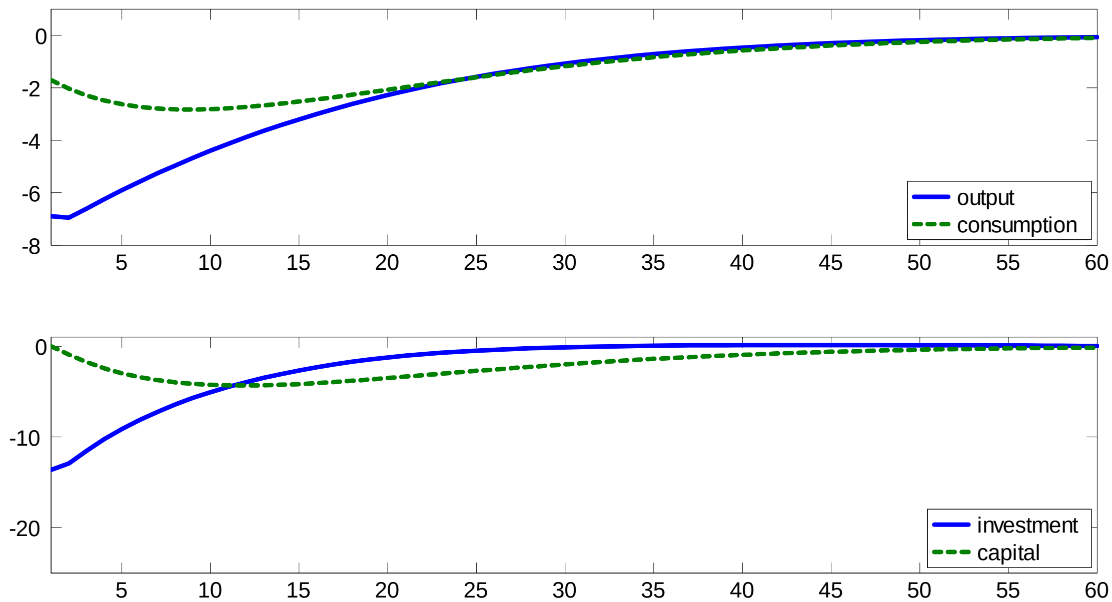

approximation which we rely upon in this version of the paper.

After an initial fall, consumption rises over time for each cohort as real interest rates increase through the recovery. The reduction in lifetime consumption is disproportionately borne by the youngest generations, the consumption of those with $j = 1$, having just begun working and with no wealth, declines by 5 percent. The older a generation, the less the reduction in their average consumption. At the date of the shock the oldest households see their consumption fall by less than 0.5 percent. This impact of the recession on the average consumption of a generation is determined by average wealth and the share of labour earnings in income. Younger generations have little wealth, and suffer more from the increase in employment risk. Older generations have more wealth to smooth out the effect on their consumption, and their lifetime income depends less on current and future labour earnings. Thus the consumption of the youngest households falls almost as much as GDP and far more than aggregate consumption. As the figure shows consumption for households of each generation over the first 5 dates of the recession, a specific value for j, on the horizontal axis, represents a different household on each cross-sectional profile. Thus we see these profiles flatten over time as households smooth the reduction in their lifetime consumption. Moreover, the very young see their consumption rise over time, while older households see it fall as it responds to the dynamics of real interest rates.

Employment risk is the primary force determining the disparate impact of the recession. Figure 7a and 7b show the third case, with no change in employment risk. The only shock

Figure 6b Consumption across generations, TFP and Employment Shocks

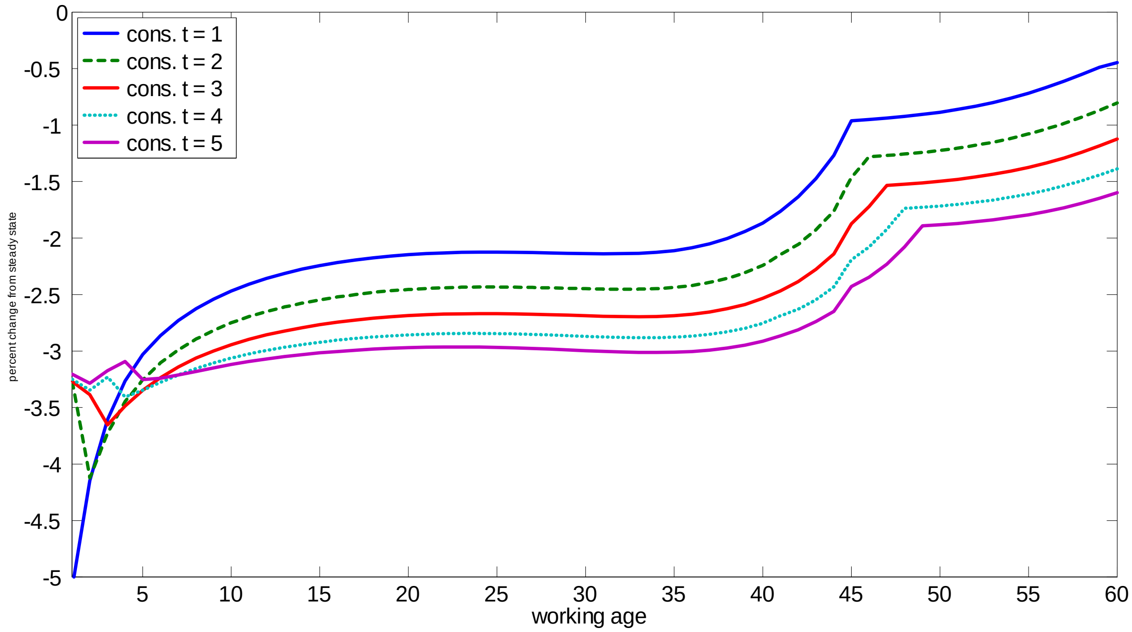

is a two percent reduction in exogenous TFP. GDP falls by the same proportion, and the fall in consumption across generations is similar.
The fall in consumption is greater for younger generations, who have less wealth to smooth the impact of the recession, but the differences across generations are small compared to the case with employment risk.

These two cases are unable to drive a large enough reduction in aggregate consumption compared to the data, where it fells over 4 percent in the recession. Starting with the shocks of the first case, in a third we add a financial shock that reduces the return on assets below the marginal product of capital. The size of the shock is specifically chosen to imply a reduction in aggregate consumption that is consistent with the magnitude observed in the 2007 recession. Further, as explained above, it has the same persistence as TFP and employment risk. In Figure 8a we see the resulting fall in GDP and consumption, investment and the aggregate capital stock.

Investment falls by 22.9 percent, compared to 13.6 in the first case without the shock to households' return on capital. In the first example, with the recession resulting from a rise in unemployment risk and a fall in TFP, the real interest rates falls to 3.38 percent, from its long-run level of 4.14 percent. In the third example that adds a shock to capital income, the real return earned by households on their wealth falls to 2.62 percent. The result is a disproportionate fall in investment, which was characteristic of the Great Recession, and a more gradual recovery in aggregate consumption. As a result, the percent deviations in

Figure 7a GDP, Cons., Inv. and Capital with TFP Shock

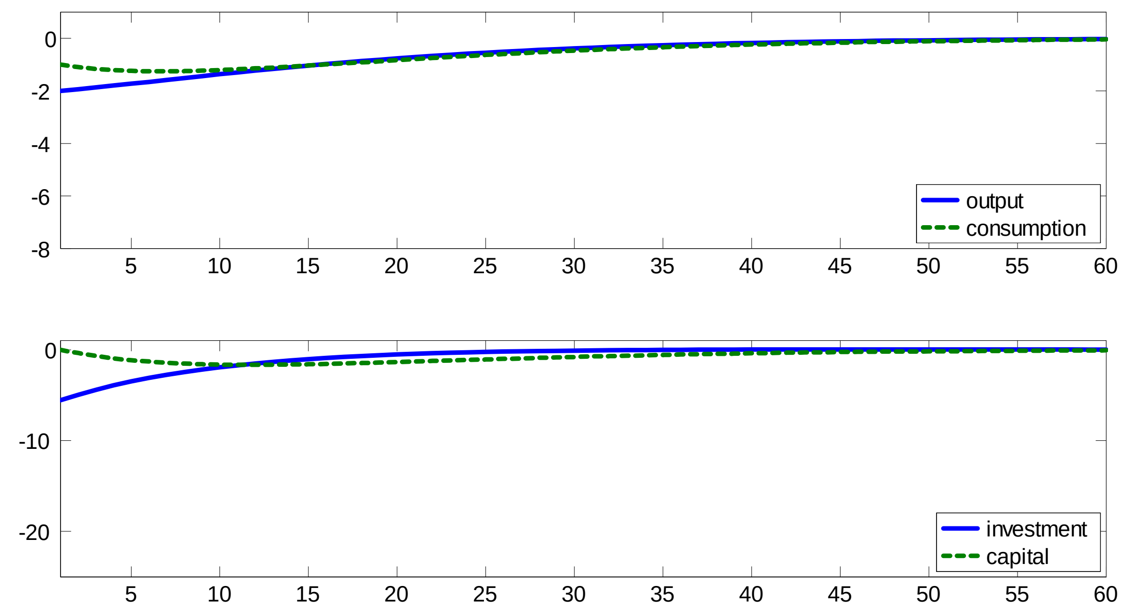

Figure 7b Consumption across generations, TFP Shock

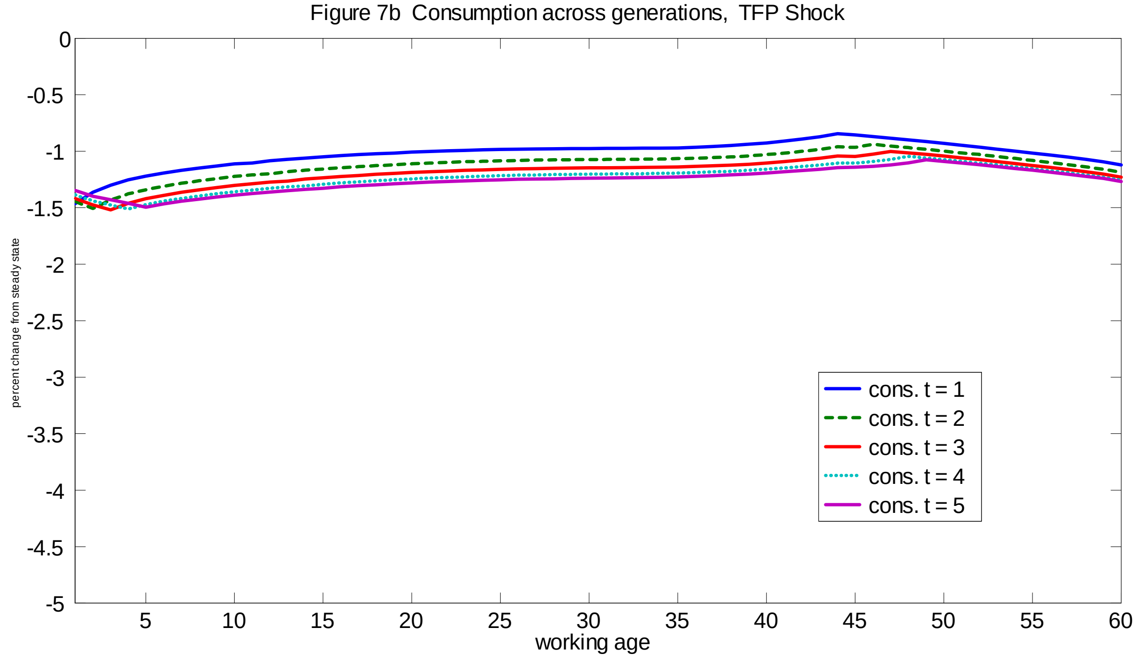

Figure 8a GDP, Cons., Inv. and Capital with TFP, Employment and Capital Shocks

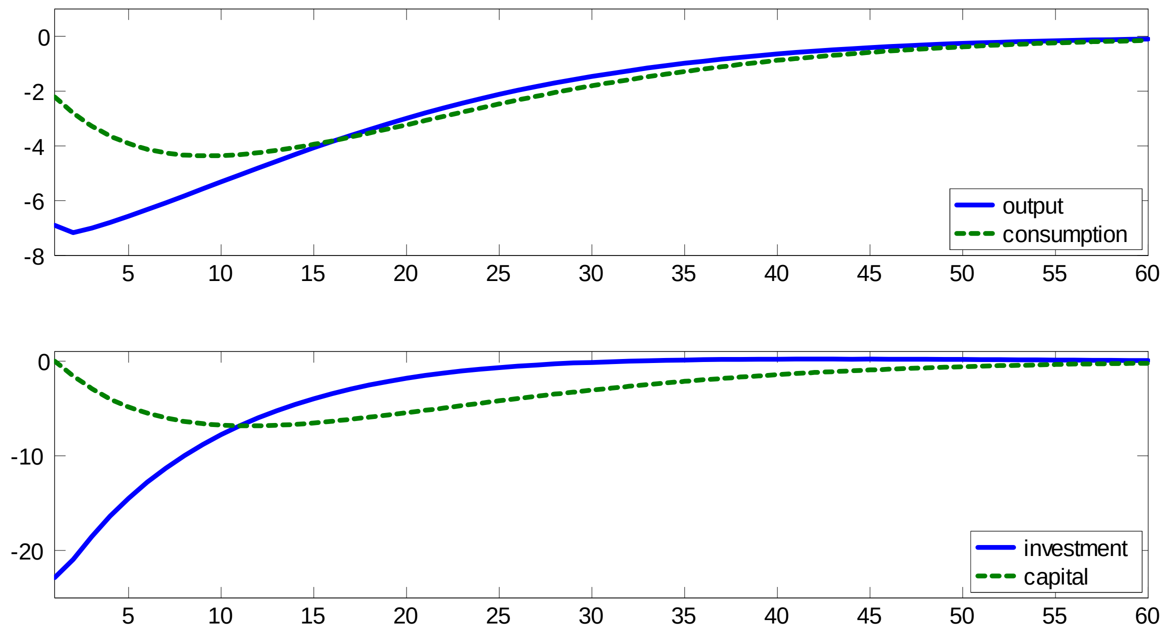

aggregate consumption, from its long-run level, eventually falls below output.

The recession now involves a larger reduction in capital income. This qualitatively changes the response in consumption across generations, and leads to larger welfare costs for the old. As apparent in Figure 8b, consumption for the oldest generations, at the start of the recession, falls by 0.8 percent compared to 0.45 percent in the first case, while the reduction in consumption for young workers is the same across the two cases. Qualitatively, the reduction in consumption is non-monotone across generations. Households that have been working between 15 and 36 years see their consumption fall by more for each additional year of work experience. These households suffer more from the loss in capital income and thus their expected lifetime wealth, than do older households who were accumulating wealth less fast before the recession.

Across these cases, we see that TFP shocks imply little redistribution across generations contemporaneously extant, and thus similar reductions in average consumption by age group over the recession. Employment and Capital Income shocks are redistributive, they imply large effects on the distribution of wealth over households of different ages and different levels of wealth. Thus with employment risk, in the first case, we see a steeper slope for the change in consumption as a function of age. The capital income shock results in a non-monotone consumption response across different ages, with the middle-aged suffering more than households that, while not new to the labour force, have less experience. The

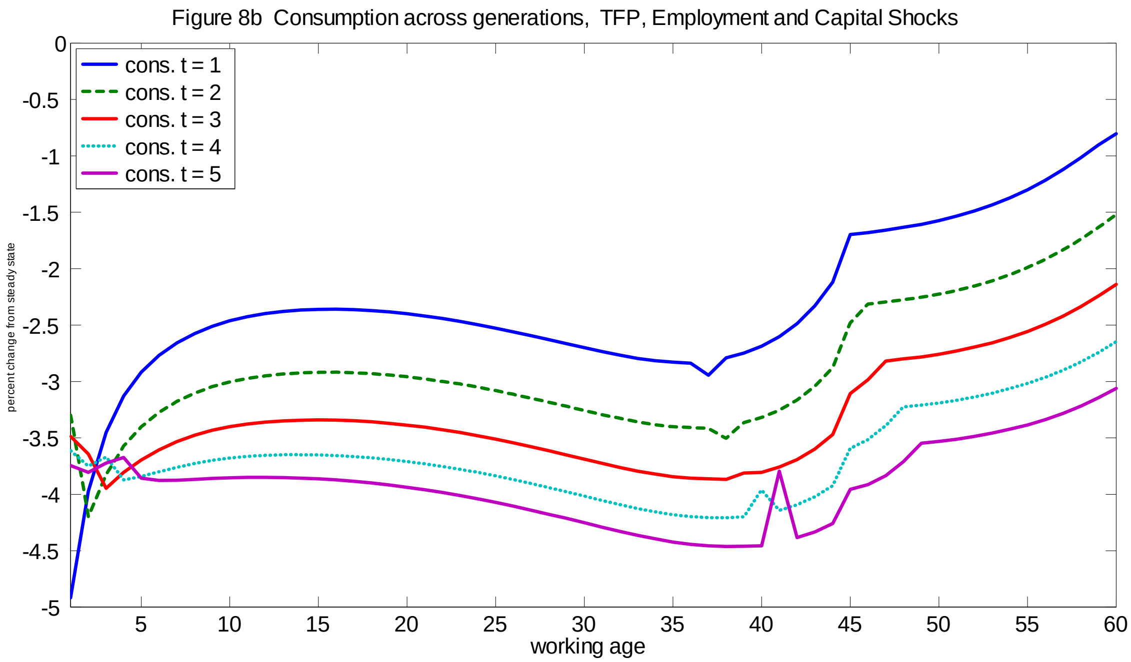

mechanics through which these shocks change the wealth of different cohorts depends on the share of capital and labour income in each generation and, moreover, their savings rates. Thus the recession has disparate effects across households, as a function of age and wealth, and the extent to which its effect differs across households depends on the shocks. Importantly, if we are to explain the fall in employment, which is impossible to explain from a frictionless representative agent model with variable labour supply (see Khan and Thomas 2013), we have to introduce a rise in unemployment risk. Explaining the large fall in consumption requires an additional shock that redistributes wealth across generations.

While aggregate shocks affects inequality in the economy, the distribution of households also determines changes in the aggregate economy. First, the capital income shock reduces aggregate consumption by distorting the consumption of middle-age households. Second, disparity in wealth determines the response in aggregate investment. If we reduce the standard deviation of income risk ten-fold, then the wealth Gini falls to 0.49 and there is less inequality within generations. Less income risk reduces households need for precautionary savings, and the aggregate capital stock falls, increasing the average real interest rate to 7 percent. Shocks to TFP, employment and capital income, in this model, reduces aggregate investment by an additional 1.7 percentage points. As households hold less wealth against income risk, they are more responsive to changes in real interest rate over the recession. This is particularly noteworthy because the percent changes in interest rates, in this case, are far less because we've not recalibrated the economy and the level of the reduction in the return on savings is the same.

5 Concluding Remarks

We have developed a quantitative overlapping generations model where households face both income and employment risk. Large levels of income risk imply that households accumulate assets to self-insure against negative shocks to labour productivity as in a standard incomplete markets model. Life-cycle savings adds considerable wealth inequality, and leads to disparate effects of shocks to employment, consistent with the fall in hours worked over the recession, and the net return on savings, required to explain the large drop in consumption. Increases in unemployment imply larger decreases in lifetime earnings and consumption for younger households with low levels of wealth, while decreases in the net return on assets have larger effect on middle-aged households with higher average levels of wealth.

High levels of income risk lead to large increases in the aggregate capital stock as households hold precautionary savings. This wealth, accumulated to offset the effects of shocks to labour productivity and employment, is relatively unresponsive to ordinary changes in real interest rates associated with TFP shocks. Thus inequality is not only affected by aggregate shocks but also shapes the aggregate response in the economy to these shocks. Levels of income risk, lower than estimated, imply more variable aggregate investment comparable, in its responsiveness to TFP shocks to a infinitely-lived household model with a representative agent. We will explore further explore the feedback from income and wealth inequality in future versions of this work, where we will allow for variable labour supply among employed households.

References

[1] Katz, Lawrence F., and David H. Autor. 1999. "Changes in the Wage Structure and Earnings Inequality." in *Handbook of Labor Economics* 3A, edited by Orley Ashenfelter and David Card. Amsterdam: North Holland.

[2] Buera, Francisco J. and Benjamin Moll (2015) 'Aggregate Implications of a Credit Crunch' *American Economic Journals: Macroeconomics* 7, no. 3, pages 1-42.

[3] Carroll, Christopher (2005) ‘The Method of Endogenous Grid Points for Solving Dynamics Optimization Problems,’ working paper.

[4] Glover, Andrew, Jonathan Heathcote, Dirk Krueger and Jose Victor Rios-Rull (2014) ‘Intergenerational Redistribution in the Great Recession,’ working paper.

[5] Guerrieri, Veronica and Guido Lorenzoni (2015) ‘Credit Crises, Precautionary Savings, and the Liquidity Trap,’ working paper.

[6] Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, Jae Song (2015) ‘What do Data on Millions of US Workers Reveal about Life-Cycle Income Risk?’ working paper.

[7] Heathcote, J., Storesletten, K. and Violante, G. L. (2010) ‘The Macroeconomic Implications of Rising Wage Inequality in the United States’ *Journal of Political Economy* vol. 118, no. 4, pages 681-722.

[8] Khan, Aubhik and Heejeong Kim (2015) ‘Segmented Asset Markets and the Distribution of Wealth,’ working paper.

[9] Kim, Heejeong (2015) ’Skill Premia, Wage Risk, and the Distribution of Wealth,’ working paper.

[10] Khan, Aubhik and Julia Thomas (2013) ‘Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity’ *Journal of Political Economy* vol. 121, no. 6, pages 1055-1107.

[11] Krueger, Dirk, Kurt Mitman and Fabrizio Perri (2015) ‘Macroeconomics and Heterogeneity, Including Inequality,’ working paper.

[12] Krueger, Dirk and Felix Kubler (2004) ‘Computing Equilibrium in OLG models with Stochastic Production’ *Journal of Economic Dynamics & Control,* 28, pages 1411 – 1436.

[13] Krusell, Per and Anthony A. Smith (1998) ‘Income and Wealth Heterogeneity in the Macroeconomy’ *Journal of Political Economy*, 106(5), pages 867 - 896.

[14] Reiter, Michael (2010) ‘Solving the Incomplete Markets Model with Aggregate Uncertainty by Backwards Induction’ *Journal of Economic Dynamics & Control* 34, pages 28–35.