Large Recessions in an Overlapping Generations with Unemployment Risk

Aubhik Khan

Ohio State University

February 2017

ABSTRACT

We explore business cycles in quantitative overlapping generations economies where households face earnings risk characterised by the high kurtosis seen in the data, and unemployment risk. The model economy exhibits large inter-generational differences in wealth driven by households' savings over the life-cycle. This substantially increases the inequality of wealth arising from households' response to uninsurable income risk, and reproduces much of the wealth inequality in the data.

Solving for equilibrium under aggregate uncertainty, we explore the mechanics of a large recession for both aggregate consumption and investment and the distribution of households. Importantly, our incomplete markets model generates declines in aggregate quantities similar to that seen in the Great Recession. As result, it allows us to evaluate the welfare costs of a large recession in an incomplete markets model consistent with the overall aggregate business cycle.

While TFP fell relatively little, in the recent recession, there was a large fall in hours worked. Consistency with these observations, in the model, requires a large, persistent increase in unemployment risk. Overall, the welfare costs of the recession vary by the share of income households derive from capital and labour. Younger workers, with relatively low levels of wealth, suffer most of the large fall in expected earnings. These welfare costs are large, given the persistence of the recession and its large impact on the earnings of finite-lived households.

Uninsurable risk and wealth inequality help shape the aggregate response of the economy. 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

Incomplete markets have seen little emphasis in modern business cycle theory. Moreover, the welfare costs of large recessions are derived using models that are inconsistent with the empirical regularities seen in aggregate quantities over the business cycle. We develop an overlapping generations model where individuals face uninsurable unemployment and, conditional on employment, earnings risk. In dynamic stochastic general equilibrium, the aggregate economy replicates the variability and comovement of GDP, consumption and investment familiar from modern business cycle theory. Further, individuals face uninsurable risk consistent with empirical estimates showing infrequent, large changes in income. The high kurtosis in individual income changes, alongside life-cycle savings by households, results in an equilibrium business cycle model with incomplete markets that reproduces much of the observed inequality of wealth seen in the data.

Several elements of the model are important in determining the distribution of households over consumption, income and wealth. First, we estimate an income shock process consistent with recent empirical estimates showing a high degree of kurtosis in individual income changes. Second, individual face uninsurable employment risk over and beyond uncertain earnings when employed. The model is calibrated to reproduce the average cross-sectional mean and median duration of unemployment over the post-war period. Further, it allows for cyclical changes in unemployment duration. Examining the effects of a large recession, we allow for a persistent rise in both mean and median unemployment duration consistent with the large increases observed following the Great Recession. A third distinguishing characteristic of our analysis, when compared to most business cycle models, are the overlapping-generations of households. Each generation starts with low average levels of wealth, savings over their working lives subject to earnings and unemployment shocks. Given borrowing limits, this makes young households, very sensitive to cyclical risk Fourth, an important determinant of the welfare costs of the business cycle across generations are Epstein-Zin preferences. These allow households to be reasonably averse to consumption risk while still fairly willing to substitute consumption over time. As a result, extensions of the model that allow for a time-varying price of capital, and persistent shocks to TFP growth, will be consistent with observed aggregate risk premia.

The recent recession involved a large fall in GDP and a persistent reduction in employment alongside a comparatively small decrease in measured TFP. In the context of the model economy, the fall in employment requires a large rise in unemployment risk. This drives aggregate changes that not only explain the observed fall in GDP but also much of the decline in aggregate consumption and business fixed investment. These aggregate changes, similar to that seen in a complete markets model arise through changes in the distribution of households over consumption, income and wealth. In the presence of borrowing limits, the fall in expected income leads younger, poorer households to sharply reducing their consumption. Older, wealthier households, holding higher levels of wealth are better able to smooth their fall in consumption relative to income. Overall, precautionary savings motives make household savings less responsive to the fall in the real interest rate over the recession. In the aggregate, there is less smoothing of consumption, which falls more, relative to investment, than in a complete markets model.

In our model economy, households income shocks follow a highly leptokurtic process consistent with recent empirical findings reported by Guvenen et al. (2015). As in a standard incomplete markets, they borrow and save in an effort to smooth their consumption against fluctuations in their income, and there is a borrowing limit. Young households are wealth-poor and, on average, have high savings rates. Households tend to accumulate wealth until retirement. These savings over the life-cycle generate substantial additional inequality, over and beyond that explained by income risk alone in an standard incomplete markets model with infinitely-lived households. As they have accumulated little savings, younger households have a higher share of income from labour, and thus experience larger lifetime effects of a persistent rise in unemployment risk on subsequent consumption. Older working households, with higher average wealth to earnings ratios, and are better able to absorb unemployment. Nonetheless, finite working lives imply larger welfare costs of persistent increases in unemployment. A persistent fall in average income, across households, leads to a permanent drop in their lifetime consumption.

Inequality of consumption, income and wealth, across and within age groups is not only affected by shocks, but also shape the aggregate response of the economy to these shocks. Realistic levels of income risk drive strong precautionary savings motive which makes households relatively unresponsive to the changes in real interest rates seen over a recession. When we reduce income risk, investment is more responsive over the business cycle.

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

[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.

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 to reproduce the fall in hours worked. The latter has strong distributional effects and drives sharp differences in the impact of the recession on different generations.

Kim (2017) studies the effect of disaster risk in a dynamic stochastic general equilibrium model with overlapping generations. She explores how changes in aggregate disaster risk affects households' precautionary savings and thus aggregate consumption. In contrast to this work, she allows for both liquid and illiquid assets. However, she abstracts from lumpy changes to individual income and thus there is far less earnings risk in her model. However, both her work and this paper are able to reproduce aggregate fluctuations familiar from complete markets models of the business cycle, for example Hansen (1985) and Prescott (1986).

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 over the recession. Thus we endogenise production and, while we do not model asset prices, we allow for capital accumulation.[2] 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 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.

We solve for stochastic equilibria in overlapping generations models with uninsurable

[2] An extension of the model, in progress, introduces capital adjustment costs and shocks to TFP growth. This setting implies large changes in asset prices and allows us to evaluate the welfare costs of business cycle, and the effect of incomplete markets on aggregate quantities, in a setting consistent with empirical risk premia.

intra-generational risk. Krueger and Kubler (2010) find that aggregate state space approximation, as developed by Krusell and Smith (1999) where moments of the distribution of capital across households are 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. Additionally, the presence of uninsurable earnings risk implies that an exact solution using Smolyak polynomials, as in Krueger and Kubler, is infeasible. As in Kim (2017), 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 age, income and wealth. This provides a solution method for overlapping generations models with many generations and uninsurable income risk.

Section 2 of the paper describes the model, while section 3 describes its the calibration, 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 save or borrow subject to a limit. A representative firm 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 $l(j)$. Each household has a persistent idiosyncratic labour productivity $\zeta_k \in \left\{\zeta_1, \ldots, \zeta_{N(\zeta)}\right\}$ that has transition matrix, $\pi_{ij} \geq 0$, $i, j = 1, \ldots, N(\zeta)$, where $\sum_{j=1}^{N(\zeta)} \pi_{i,j} = 1$ for $i = 1, \ldots, N(\zeta)$.

Households face a risk of unemployment that is uncorrelated with their labour productivity. Let $e \in \left\{e_1, \ldots, e_{N(e)}\right\}$ describe employment. We will assume $e_1 = 1$ and $e_{N(e)} = 0$, corresponding to full employment and unemployment over a period. Intermediate values correspond to partial unemployment, of duration less than one model period. The probability over employment is $\pi_i^e(s) \geq 0$, $i = 1, \ldots, N(e)$ with $\sum_{i=1}^{N(e)} \pi_i^e(s) = 1$. 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). However, in contrast to that paper, the labour force is not only determined by the aggregate exogenous state but also varies with the existing distribution of employment across households. 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

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

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 $(1-\tau)\, w\zeta l\,(j)\, e$ 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. Assuming social security benefits are proportional to labour income if the household were still working, given the infrequent nature of changes to income process we will use, allows them to vary with lifetime earnings.

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

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

We allow for non-expected utility and, following Epstein and Zin (1989), assume $\sigma > 0$ as the inverse of the elasticity of intertemporal substitution and $\gamma > 0$ as the coefficient of relative risk aversion. 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, $\{\pi^s_{mn}\}^{N(s)}_{m,n=1}$ 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.

$$
v\left(J,a,e,\zeta_k,s_m,A\right)=\left(\left(1-\beta\right)\left(\left(1+r\right)a+b\zeta_k\right)^{1-\sigma}\right)^{\frac{1}{1-\sigma}}. \qquad (2)
$$

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.

$$
\begin{aligned}
v\left(j,a,e_h,\zeta_k,s_m,A\right) &= \max_{c,a'}\Bigg(\left(1-\beta\right)c^{1-\sigma} && (3)\\
&\quad +\beta\Big(\sum_{m=1}^{N(s)}\pi^s_{mn}\sum_{i=1}^{N(e)}\pi^e_i\left(s_n\right)\sum_{l=1}^{N(\zeta)}\pi_{kl}v\left(j+1,a',e_i,\zeta_l,s_n,A'\right)^{1-\gamma}\Big)^{\frac{1-\sigma}{1-\gamma}}\Bigg)^{\frac{1}{1-\sigma}}\\
&\text{subject to}\\
c+a' &\leq (1+r)\,a+b\zeta_k\\
c &\geq 0,\ a'\geq \underline{a}.
\end{aligned}
$$

Retired households' income do not face employment risk, but we allow their pensions to vary with labour productivity, ζ. An alternative, which we have also explored, is to keep

ζ constant at its level when last working. This makes little difference to the results, but While there is no employment, we include e in their value functions for consistency with (1) when $j = J_r - 1$, the only risk they face is aggregate.

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. The aggregate labour supply is the sum of working age households, over labour productivity, that are employed,

$$L = \sum_{j=1}^{J_r-1} \int_{\left\{e_1,\ldots,e_{N(e)}\right\}\times\left\{\zeta_1,\ldots,\zeta_{N(\zeta)}\right\}\times[\underline{a},\infty)} e\zeta\mu\left(j, d\left[e \times \zeta \times a\right]\right).$$

Let $\mu(j, e, \zeta, a)$ describe the measure of households of working age j with labour productivity ζ, employment status, e and wealth a. The aggregate capital stock is

$$K = \sum_{j=1}^{J_r} \int_{\left\{e_1,\ldots,e_{N(e)}\right\}\times\left\{\zeta_1,\ldots,\zeta_{N(\zeta)}\right\}\times[\underline{a},\infty)} a\mu\left(j, d\left[e \times \zeta \times a\right]\right).$$

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

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

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

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

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_h^e and π_k^0. The probability distribution $\{\pi_k^0\}_{k=1}^{N(\zeta)}$ describes the invariant distribution of households over labour productivity, determined by $\{\pi_{kl}\}_{k,l=1}^{N(\zeta)}$. Lastly, let $\left\{\mu_j^a\right\}_{j=1}^{J}$ describe the number of households in each cohort $j = 1, \ldots, J$. In the following, we assume that $\mu_j^a = \frac{1}{J}$ for all j.

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

$$\mu(j+1, e_h, \zeta_k, B) = \pi_h \sum_{i=1}^{N(\zeta)} \pi_{ik} \int_{\{(e,\zeta,a)|g(j,e,\zeta,a)\in B\}} a\mu(j, d[e \times \zeta \times a])$$

where $g(j, e, \zeta, a)$ is the savings decision rule for a household of age j with current employment status e, labour productivity ζ and beginning of period wealth, a.[3] The distribution $\mu(j, e, \zeta, a)$ is the endogenous component of the aggregate state, A. We study its evolution in recursive general equilibrium with aggregate shocks.

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 $J_r = 40$. In developing a framework that allows for realistic risk premia, we allow for a fairly high elasticity of intertemporal substitution and set $\sigma = \frac{2}{3}$ alongside a coefficient of relative risk aversion, γ, of 5. These values go a long way in reasonable equity premia and a risk-free real interest rate in a complete market model with long-run TFP risk and capital adjustment costs. As I intend to integrate these elements into the current model, I use preferences consistent with risk premia.

The subjective discount factor, $\beta = 0.93$ implying, in our overlapping generations model, a 4 percent average 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 43.5 percent of annual earnings, $b_0 = 0.435w$. The social security replacement rate is $b = 0.45$.

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

3.1 Earnings Shocks Estimation

We estimate a leptokurtic income shocks process following Kaplan, Moll and Violante and contrast it to a more log-normal conventional income process with transitory and persistent components, taken from Khan and Kim (2014). I describe first the benchmark income process that reproduces the high degree of kurtosis apparent in the moments reported from male earnings in Social Security Administration data by Guvenen et al. (2015).

Table 1: Individual Earnings Process

moment	data	model
variance annual log earnings	0.70	0.74
variance 1 year change	0.23	0.32
variance 5 year change	0.46	0.53
kurtosis 1 year change	17.8	17.7
kurtosis 5 year change	11.6	11.7
fraction 1 year change < 10%	0.54	0.59
fraction 1 year change < 20%	0.71	0.69
fraction 1 year change < 50%	0.86	0.90

We assume an income process that is similar to that estimated by Kaplan et al (2016), the parameters are estimated using simulated methods of moments to minimise, for our annual calibration, the sum of squared residuals with the eight empirical moments above.

$$
\begin{aligned}
\log \zeta_t &= \zeta_{1,t} + \zeta_{2,t} \\
\zeta_{1,t+1} &= \rho_1 \zeta_{1,t} + \chi_{1,t} \varepsilon_{1,t} \\
\zeta_{2,t+1} &= \rho_2 \zeta_{2,t} + \chi_{2,t} \varepsilon_{2,t}
\end{aligned}
$$

Both ε_t^1 and ε_t^2 are normally distributed innovations with means of 0 and variances of σ_1^2 and σ_2^2, respectively. The random variables $\chi_{1,t}$ and $\chi_{2,t}$ take on values of 0 or 1 with Poisson Probabilities λ_1 and λ_2. When $\chi_{i,t} = 1$, there is an innovation to the stochastic process $\zeta_{i,t}$, $i = 1, 2$, otherwise $\zeta_{i,t+1} = \rho_i \zeta_{i,t}$, $i = 1, 2$. Notice that $\zeta_{1,t}$ and $\zeta_{2,t}$ are uncorrelated, and each is subject to infrequent shocks when $\lambda_i \in (0, 1)$. Low values of λ_i with a large variance for ε_i and ρ_i imply a very leptokurtic shock process.

Simulating income for our annual model, we estimate the 6 parameters using a million period simulation. The resulting match to the data is shown above in table 1 and the estimated parameters values are reported below.

Table 2: Income Process Parameter Estimates

parameter	ρ_1	ρ_2	σ_1^2	σ_2^2	λ_1	λ_2
value	0	0.95	2.6716	0.3428	0.0183	0.500

We contrast this income process with a conventional shock process, without excess kurtosis, 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.[4]

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

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.

[4]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.

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, $\widetilde{v}_{i,j,t}$.[5] 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 $\log \epsilon^v_{i,j,t} \sim N(0, \sigma^2_{v_t})$.

Following Heathcote, Storesletten, and Violante (2010), we estimate year-varying shock variances $\{\sigma^2_{\eta_t}, \sigma^2_{v_t}\}$, 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. Calibrating to 2004, we use 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.

Beyond shocks to earnings for employed workers, the model is consistent with the average unemployment rate and mean and median duration of unemployment. In choosing the average values of the mean and median duration of unemployment, average monthly data between two periods. The first precedes the Great Recession, and is from December 1981 to December 2007. The second is from January 2008 to April 2016. Over the first period, the mean duration of unemployment was 15.9 weeks and the median duration was 7.8 weeks. In the second period, the mean duration of unemployment was 31.6 weeks and the median was 16 weeks. It is useful to notice that, while the trough of the recession was in the second quarter of 2009, unemployment duration rose slowly after the recession,

[5] 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.

peaking at the end of 2011. Averaging these two samples, we target a mean duration of unemployment of 0.4567 fraction of a year, and a median of 0.2286.

We set $N(e) = 3$ with a support of $\{1, 0.7713, 0\}$. Next we set $\frac{p_2(1-e_2)}{p_2+p_3} = 0.4567$ and $p_2 + p_3 = 0.05$. This reproduces the mean and median durations of unemployment as well as an overall unemployment rate of 5 percent. We also allow p_2 and p_2 to vary with s, the exogenous aggregate state. Assuming aggregate TFP takes on 3 values, in a recession we raise $p_2 + p_3$ to imply a 10 percent unemployment rate, then adjust p_2 to reproduce the mean duration of unemployment observed over the second sample period. Similarly, in an expansion, we choose p_2 and p_2 to imply a rise in total hours worked that is equal to the fall seen in a recession. Finally, we set p_2 to reproduce unemployment duration between 1981 and 2007.

As we target both mean and median duration of unemployment, allowing these probabilities to evolve with TFP, aggregate hours worked, in the model, lags TFP. Thus the unemployment rate does not immediately adjust to the long-run level consistent with TFP as in Krusell and Smith (1998). We estimate a log-normal AR(1) process for TFP data on the aggregate capital stock, derived from the fixed asset tables, and total hours worked data provided by Cociuba et al. (2012) and extended through 2012Q4. The persistence of annual total factor productivity is 0.909 and the standard deviation of innovations is 0.01447.

3.2 Solution

We discretise individual earnings and solve decision rules The leptokurtic income shocks process estimated to match the moments from the SSA data is discretised using a 25 million period simulation. We choose 12 points in its support and the simulation values from the estimated continuous shock process onto this discrete support. The frequency of changes from one discrete grid value to another provides the Markov Chain. The support is adjusted for efficiency, with grid points space far apart near the endpoints, and more narrowly concentrated near the centre. When we use the log-normal shock process for income estimated using the PSID, the Rouwenhorst algorithm is used to combine the transitory shock, using 2 grid points, and the persistent shock, using 5 grid points.

Household value functions are solved using a generalisation of the Backwards Induction algorithm of Reiter (2010) described by Kim (2017) who applies it to solve for dynamic stochastic general equilibrium of an overlapping generations model with two assets that vary in their liquidity. Reiter's method abstract from the simulation step of the Krusell and Smith (1999) method of approximate state aggregation. While it uses aggregate state space approximation, it solves for the law of motion of this aggregate state with individual decision rules. This involves a proxy distribution, at each value of the approximate aggregate state, which allows for the computation of next period's aggregate state when we solve household decision rules.

We use the aggregate capital stock as our approximate aggregate state. At each aggregate grid point (s,k), given a proxy distribution $\mu(j, e, \zeta, a; s, k)$, we conjecture a value for $\widehat{k}'$, the future approximate aggregate state. Next, determine the actual end of period aggregate capital stock by solving household value functions at each age, labour productivity and employment level. While households have finite lifetimes, this involves value function iteration over aggregate states as we iterate on the conjecture $\widehat{k}'$ until it's consistent with the actual savings of households.

Starting with $j = J$, we determine $v^{n+1}(J, e, \zeta, a, s, k)$ using (**??**) where n indexes the iteration of value functions, across ages, with respect to the aggregate state (s, k). Using $v^n(j+1, e', \zeta', a', k')$, $j = 1,..., J-1$, we solve decision rules using the endogenous grid method of Carroll (2005). Alongside $v^{n+1}(j, e, \zeta, a, k')$ this provides decision rules $a' = g^{n+1}(j, e, \zeta, a, k)$ and thus, using the proxy distribution, and an actual end of period capital stock, k'. We iterate until $\widehat{k}' = k'$.

The accuracy of this method relies on choosing a set of proxy distributions that are representative of the distribution of households over a simulation. Using the decision rules derived from backwards induction, we simulate the model for 1100 periods, discarding the first 100 observations. In contrast to the method of Krusell and Smith, this simulation is not necessary to update the aggregate law of motion for the approximate aggregate state. As its only purpose is to refine the proxy distributions, the necessary length of simulation, the slow step in the Krusell and Smith algorithm when applied to models with multi-dimensional heterogeneity, is far shorter. The simulation provides reference distributions. These are mapped onto a new set of proxy distributions. However, as the reference distributions are not moment consistent, their mean does not equal the value of the approximate aggregate state at any grid point, we must derive proxy distributions by solving a quadratic minimisation problem than constrains the proxy distributions to be consistent with the conditional distribution over (e, ζ), at each reference distribution, and to have the same mean as the approximate aggregate state. This problem is too large to be tractable without state space reduction of the distribution. The innovation in Kim (2017) provides a feasible reduction, aggregating over (j, e, ζ) using weights $\omega_{j,e,\zeta}(a)$, solving for a reduced dimensional proxy distribution, then recovering the full distribution. The details are in her paper, and this additional state state reduction, over and beyond the use of an approximate aggregate state, makes the application of Backwards Induction feasible to solve for stochastic equilibria in overlapping generations models with multi-dimensional distributions. The method is very accurate, as well as fast. Den Haan's accuracy measure gives a maximum error of 1.14 percent which, as reported by Khan and Thomas (2013) is comparable to that of the stochastic beta model of Krusell and Smith.

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 implied by the Markov Chain of leptokurtic income shock process. Over time they borrow or save, and their earnings grow with experience. Over their life-cycles, on average 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 of inequality is the result of differences in earnings shocks.

Table 3: 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	0.14	0.37	0.49	0.90	1.0	0.62	0.05

Table 3 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 15 percentage points of the SCF data in 2004. Nonetheless, in the absence of a motive to accumulate wealth over and beyond consumption smoothing and retirement implies less skewness in its distribution than in the data. Thus the wealthiest 1 percent of households hold 14 percent of the total wealth of the economy, while, in the data, the corresponding number is 29 percent. However, as seen in figure 1 the overlapping generations economy, with an estimated income shock process, is able to generate an unusually large level of inequality.

Figure 2 shows the distribution of wealth for households of all ages. The right horizontal axis is the logarithm of earnings, while the front axis shows wealth. Most households have

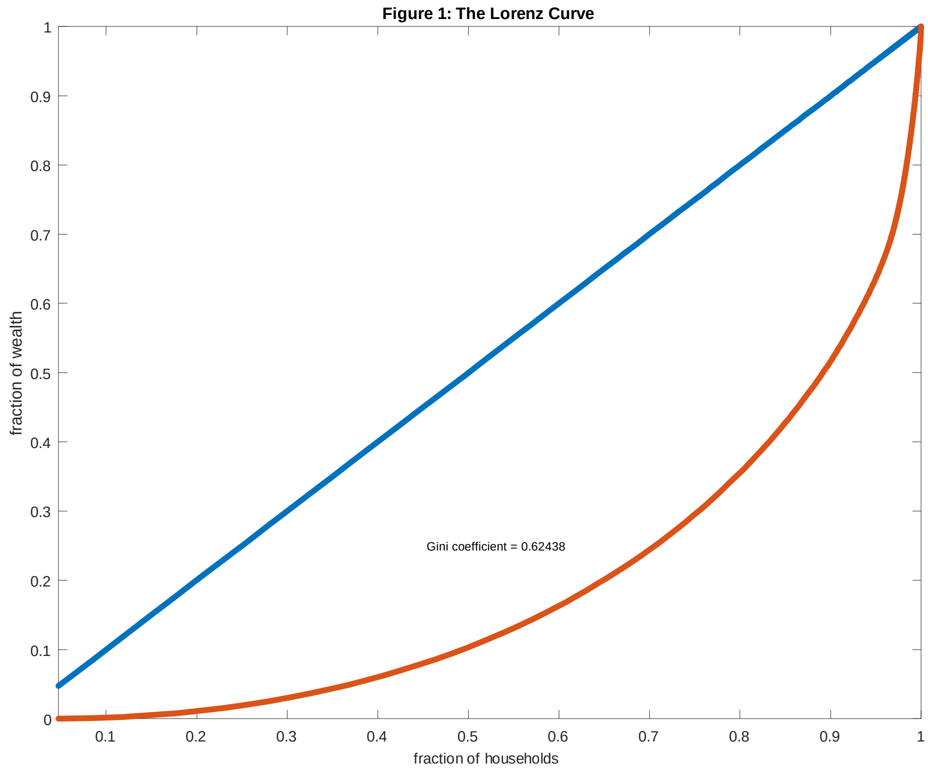

relatively little wealth and the distribution of wealth across households is very skewed. There is considerable dispersion in the distribution of wealth, and higher income households tend to be wealthier given the persistence in the estimated income shock process.

Interestingly, the leptokurtic income shock process actually increases inequality. If we use the alternative estimated log-normal shock process, there is less concentration of wealth. The top 1 percent of households then hold 8.8 percent of the aggregate capital stock. However, the Gini coefficient rises to 64 percent. The accumulation of wealth over working lives, towards financing consumption after retirement, drives considerable inequality even in the absence of income risk. If we reduce the standard deviation of these log-normal earnings shocks 10 times, inter-generational wealth inequality remains high and the Gini coefficient is still 0.46.

Figure 3 shows the average level of income, consumption and wealth in a generation over time. Its members 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. The low probability of a large, transitory income shock results leads to fluctuations in average income in our sample of 100,000 households, in the top panel. As mentioned in the beginning of section 3, the social security replacement rate used in the

Figure 2 The Distribution of Wealth

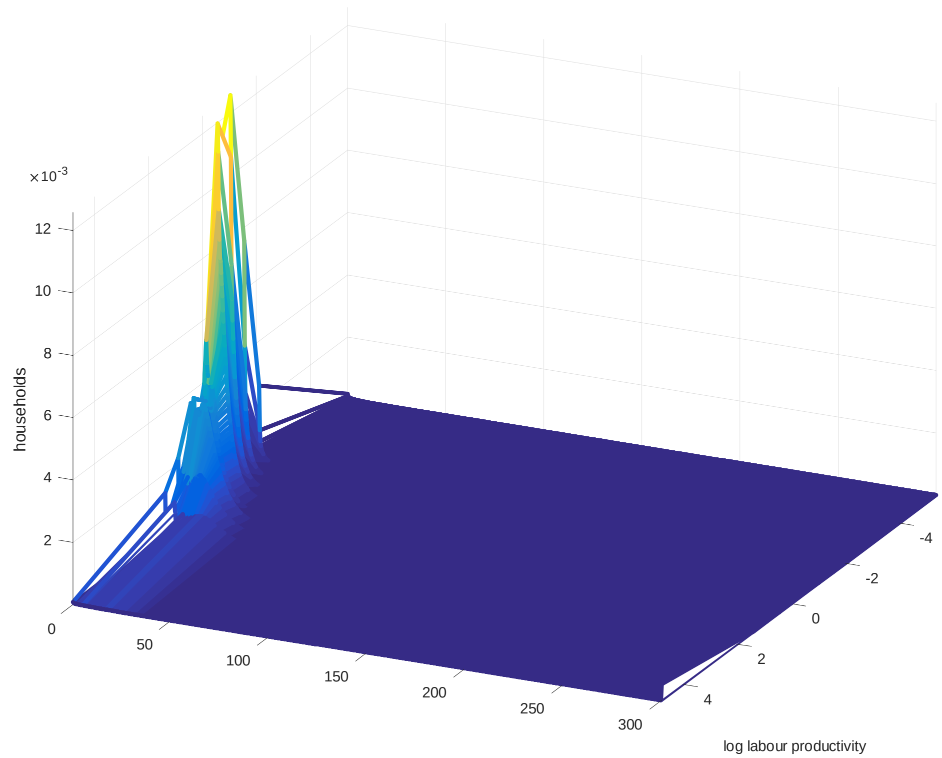

model is 0.45, and retirement earnings are proportional to labour income. This income, averaged across households, is shown in the top panel starting with $j = 45$.

The equilibrium real interest rate in the steady state is 4.09 percent, while household's subjective discount factor is $\beta = 0.93$. Income risk and life-cycle savings drive the accumulation of wealth seen seen in the bottom panel. On average, households save until retirement, then gradually decummulate their assets to smooth consumption relative to income. While working, the average rise in income and wealth implies a gradual increase in consumption. Consumption initially falls in retirement, as households' marginal rate of substitution is less than one. The rise in consumption at later is the result of a decline in the precautionary savings motive which, even in retirement, is used to smooth consumption against unexpected changes in pensions.

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

Figure 3 A Cohort of 100000 Households

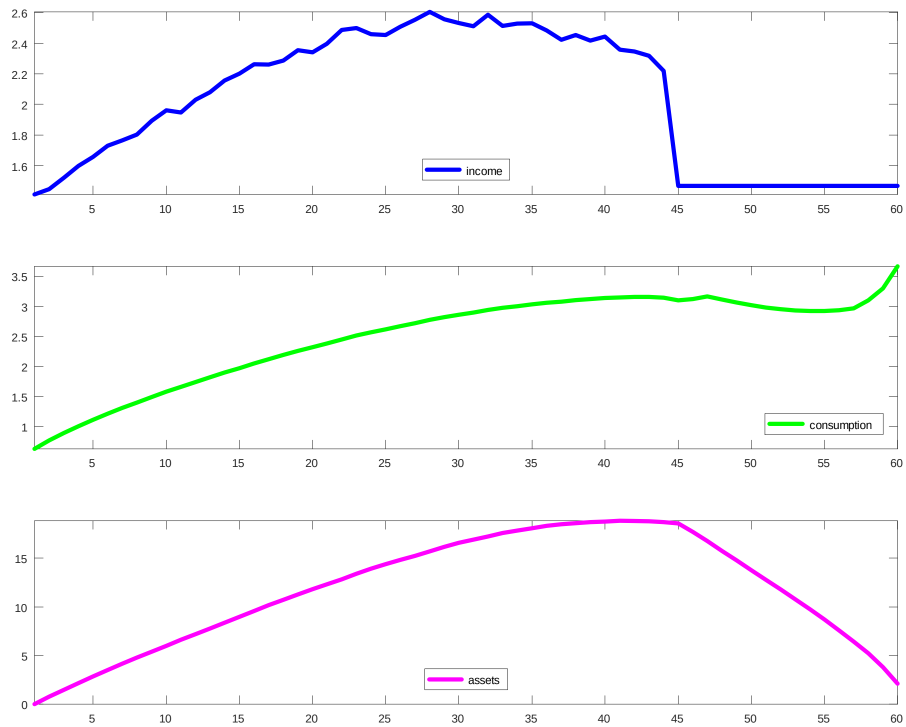

4.2 Business Cycles

Table 4 reports business cycle moments from a 1000 period simulation of the benchmark model where households face uninsurable earnings and employment risk. Aggregate fluctuations are driven by exogenous shocks to total factor productivity and the distribution of employment risk. As the probability of unemployment rises exogenously in recessions, and falls in expansions, the number of employed workers varies over the business cycle.

Table 4: Business Cycle Moments[6]

	Z	Y	C	I	N	K	r	w
std	1.30	1.97	0.67	5.31	1.51	0.69	0.22	1.11
correlation	0.86	1.00	0.75	0.93	0.83	0.20	0.94	0.65

The incomplete markets model is able to reproduce the typical business cycle at least as well as a complete markets equilibrium business cycle model. There are several notable

[6] All series are HP-filtered with a weight of 100. The series are TFP (Z), GDP (Y), aggregate consumption (C), aggregate investment (I), labour (N), capital (K), the real interest rate (r) and the real wage (w). The first row reports percent standard deviations while the second reports the contemporaneous correlation of each series with output.

departures from the latter. First, incomplete markets with a strong precautionary savings motive, implied by the high risk aversion characterising our Epstein-Zin preferences, reduce the correlation between consumption and GDP when compared to that seen in a model with iso-elastic utility. If we compare the above results to a model with expected utility and the same elasticity of intertemporal substitution, the correlation rises to 0.93. In the benchmark model, households hold fairly high levels of savings to smooth their consumption against the possibility of large, infrequent shocks to earnings. This makes consumption less responsive to aggregate shocks which, in turn, drives a larger percent variability in investment. Second, exogenous changes in employment risk lead to larger movements in total hours worked than seen in a model with variable labour supply. Further as unemployment risk slowly propagates through the distribution of working-age households, employment lags GDP as it does in the data. In Table 4 this is seen in a lower contemporaneous correlation between Y and N. Employment lags output, and the correlation with lagged GDP is 0.85. In contrast, the correlation between Y_{t-1} and N_t is low, at 0.34. While unemployment is exogenous, as in the model of Krusell and Smith, its implications for aggregate employment are more consistent with the data than that seen in a model with endogenous labour supply. The contemporaneous correlation between the equilibrium real wage and GDP is 0.65, lower than that usually seen in a complete markets equilibrium business cycle model with variable labour supply. Both the lower correlation between consumption and GDP, and the lag in employment, move the incomplete markets model business cycle closer to the data when compared to the standard expected utility representative agent model.

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.[7] Surprisingly, measured TFP fell by only 2.18 percent over the same period. We explore the impact, on the model economy, of shocks to empirically consistent shocks to TFP and employment.

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 a persistent shock to unemployment risk. This raises the unemployment rate to 10 percent from its ordinary level just above 5 percent. Figure 4 plots the persistent shocks to TFP and employment.

The model is able to explain the decline in GDP and much of the fall in investment observed in the 2007 recession. As shown in figure 5, GDP fall 5.3 percent upon impact, and investment drops 13.1 percent. Aggregate consumption declines gradually through

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

Figure 4 TFP and Employment Shocks

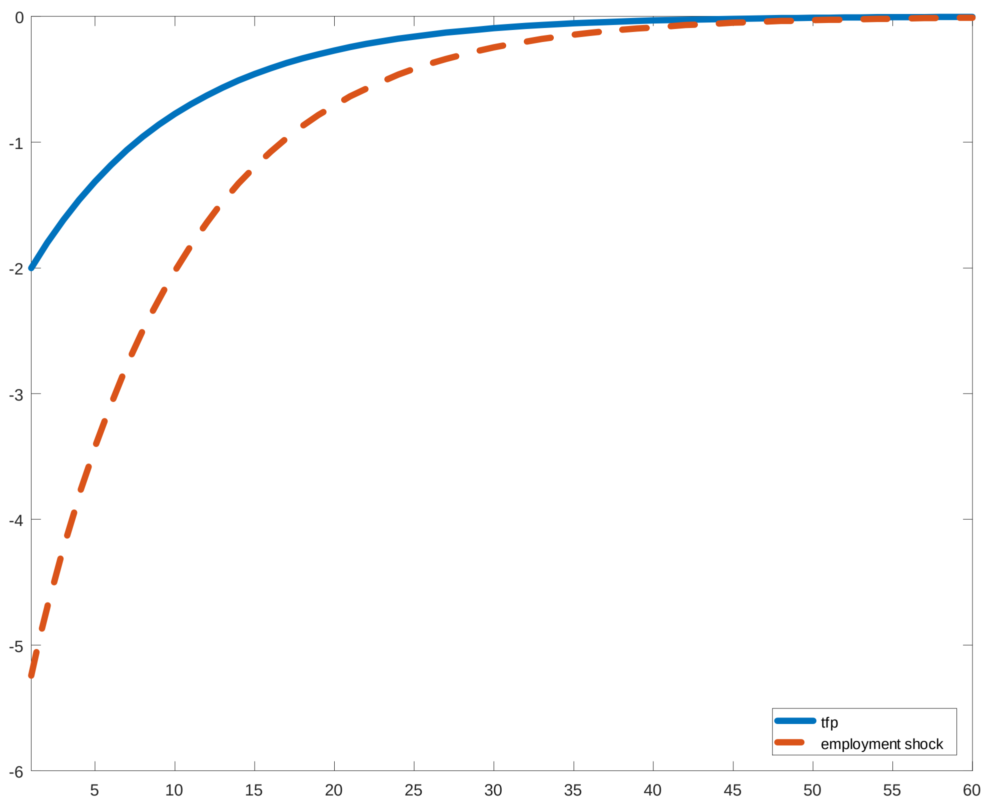

Figure 5 GDP, Cons., Inv. and Capital

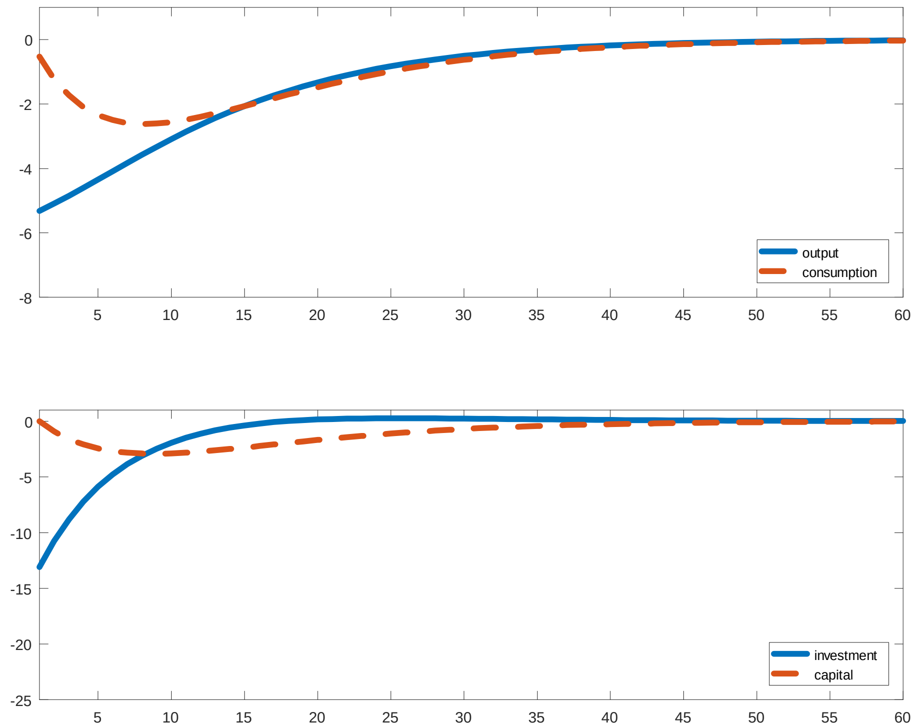

a familiar hump-shaped response. The aggregate shocks to the economy imply a large negative wealth effect, especially for working age households who experience both a fall in their earnings, the probability of employment and the return to savings. However the initial fall in consumption is dampened by a substitution effect that reduces savings over and beyond that implied by consumption smoothing. This is in response to the fall in the expected real interest rate, which reduces the relative price of current consumption. Over time, as TFP begins to recover, the rise in the real interest rate increases savings rates and aggregate consumption falls further. Eventually, 8 periods after the shock, it declines 2.61 percent. This is 64 percent of the fall seen in the Great Recession.

The welfare effects of the recession, driven by shocks to total factor productivity and employment risk, vary across households. This is evident in Figure 6a 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, the figure offers insight into the short term impact of the recession across cohorts.

In Figure 6b, I compute the average welfare loss, in units of steady state consumption, for the distribution of households over wealth and income in each cohort. The steady

Figure 6a Consumption across generations

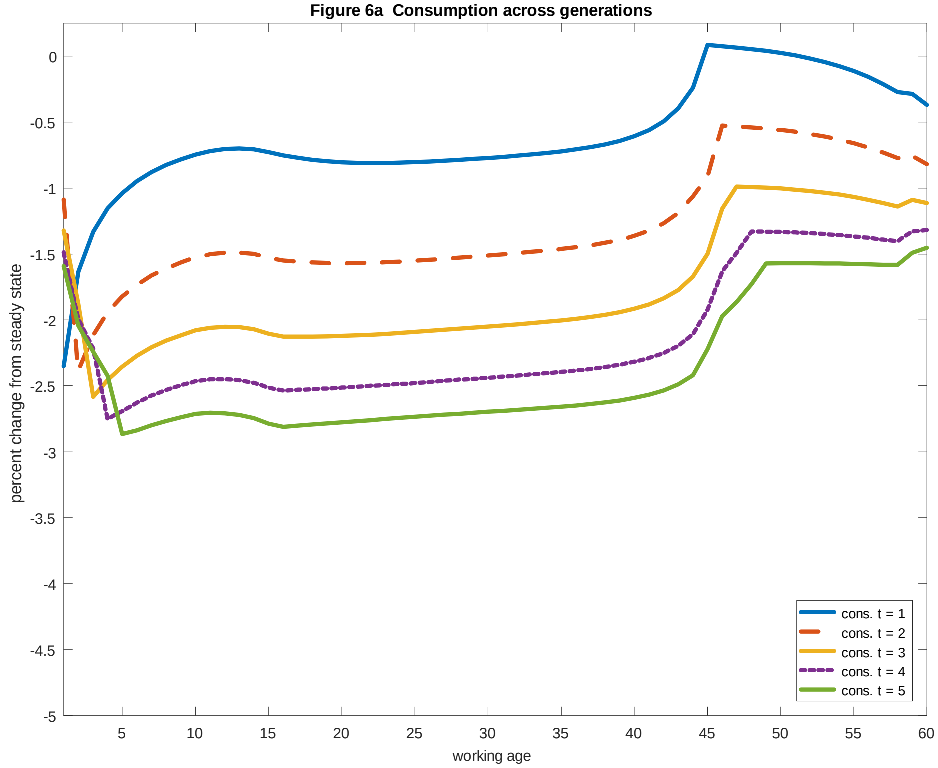

state distribution is used to weight households of each age, at the start of the recession. A common percent loss in consumption, that equates average utility to the level after the TFP and employment shock, is reported. Comparing these two figures, we see that the initial fall in consumption, for the youngest households is more severe than their overall loss in lifetime welfare. However the impact of the recession falls disproportionately on younger households, and it is less for older households. The discontinuity in the welfare measure occurs at the retirement age where households become impervious to higher employment risk in the recession. As this employment risk increases the expected loss of labour income, over and beyond that which arises from the fall in real wages, we see that younger households that are more dependent on earnings experience a larger fall in their expected lifetime utility.

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 almost 2.5 percent. across working ages, the older a household the less the reduction in its consumption. At the date of the shock the oldest retired 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

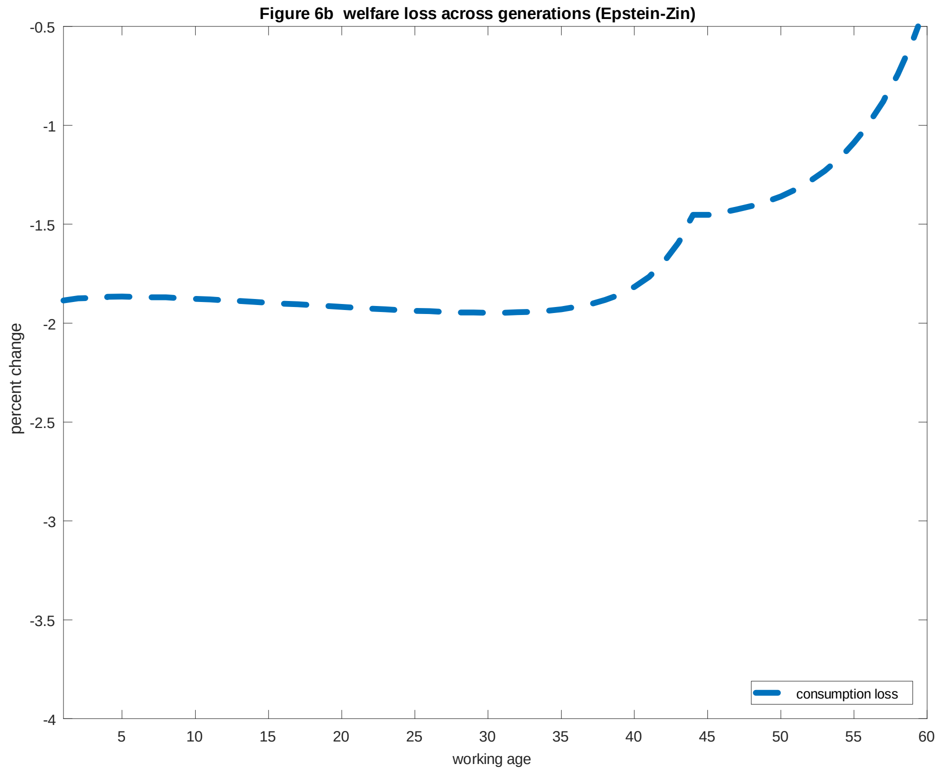

the share of labour earnings in income. Younger generations have little wealth, and suffer more from the increase in employment risk. Older generations have more wealth to smooth out the effect on their consumption, and their lifetime income depends less on current and future labour earnings. Thus the consumption of the youngest households falls 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.

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. This results in a larger welfare cost of the recession on younger households. If we eliminate the rise in unemployment, figure 6b becomes flatter, and there is less difference in the incidence of the recession over age.

While aggregate shocks affects inequality in the economy, the distribution of households also determines changes in the aggregate economy. 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 0.47 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 8.3 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 earnings as in a standard incomplete markets model. This precautionary savings motive is strengthened by a high aversion to risk, and a leptokurtic distribution of earnings shocks. In addition, 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. Increases in unemployment imply larger decreases in lifetime earnings and consumption for younger households with low 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.

The present work is a progress report towards a general study of the effect of idiosyncratic and aggregate risk in incomplete markets models. An extension of the model described here allows for changes in the price of capital and persistent shocks to the rate of growth of TFP. In a complete markets setting, these elements imply a realistic return on equity, and a low risk-free real interest rate. The implied empirically plausible risk premia, in the current setting, will allow us to better understand the effect of changes in unemployment risk, as well as aggregate TFP risk, for the business in an incomplete markets model.

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] Cociuba, Simona E., Edward C. Prescott, and Alexander Ueberfeldt (2012) "U.S. Hours and Productivity Behavior Using CPS Hours Worked Data: 1947-III to 2011." Manuscript, Dept. Econ., Univ. Western Ontario working paper.

[5] Epstein, Larry and Stanley Zin (1991) 'Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis' *Journal of Political Economy*, vol. 99, pages 263-286.

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

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

[8] 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.

[9] Hansen, Gary (1985) 'Indivisible Labor and the Business Cycle,' *Journal of Monetary Economics*, vol. 16, pages 309–327.

[10] 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.

[11] Kaplan, Gregory, Benjamin Moll and Giovanni Violante (2016) 'Monetary Policy According to HANK' working paper.

[12] Kim, Heejeong (2017) 'Inequality, Portfolio Choice and the Business Cycle' working paper.

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

[14] 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.

[15] Krueger, Dirk, Kurt Mitman and Fabrizio Perri (2017) '*Macroeconomics and Household Heterogeneity*,' in Taylor, John and Harald Uhlig (eds.) Handbook of Macroeconomics, vol. 2, pages 843-921, Elsevier.

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

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

[18] Prescott, Edward C. (1986) 'Theory Ahead of Business Cycle Measurement,' Federal Reserve Bank of Minneapolis *Quarterly Review*, vol. 10, pages 9–22.

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