Growth, Uncertainty and Business Cycles in an Overlapping Generations Economy

Aubhik Khan

Ohio State University

Ben Lidofsky[1]

Ohio State University

February 2019

ABSTRACT

We explore business cycles in overlapping generations economies driven by shocks to total factor productivity growth. Households have uncertain lifetimes and face both uninsurable earnings and employment risk. Unemployment risk is countercyclical and we allow for time-varying volatility of aggregate shocks. Households, with non-separable preferences, save using capital and the relative price of capital varies over time. This setting–with both idiosyncratic and aggregate uncertainty shocks–drives large, countercyclical risk premium and generates high levels of precautionary savings. We find that changes in precautionary savings have important implications for aggregate consumption.

Persistent negative shocks to TFP growth, and increases in their uncertainty, drive large declines in consumption. This helps explain the slowdown observed since the onset of the Great Recession. An empirically consistent, moderate shock to TFP growth rates implies a large and persistent fall, against trend, in aggregate consumption. Moreover, a rise in aggregate uncertainty reconciles the recovery in TFP growth rates with a more protracted decline in consumption.

Uninsurable income risk helps shape the aggregate response of the economy over the business cycle. Changes in households' precautionary savings motives not only affect the distribution of wealth, it also changes the volatility of aggregate consumption and investment.

[1]Aubhik Khan: mail@aubhik-khan.net; Ben Lidofsky: benjamin.lidofsky@gmail.com.

1 Introduction

We develop an overlapping generations model where individuals face uninsurable unemployment and earnings risk. Aggregate technology is subject to trend-stationary shocks to TFP growth with time-varying volatility. Thus, we study business cycles driven by both growth and uncertainty shocks. In dynamic stochastic general equilibrium, the aggregate economy replicates the variability and comovement of GDP, consumption and investment familiar from modern business cycle theory.

Uninsurable income risk, 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. Moreover, this distribution of households over income and wealth has aggregate implications. When households face counteryclical income risk, and hold large levels of precautionary savings, aggregate consumption is actually less responsive to exogenous changes in TFP growth. Nonetheless empirically consistent aggregate shocks drive large, persistent declines in consumption and investment comparable to those seen in the data.

Several elements of the model are important in determining the distribution of households over income, wealth and consumption. First, idiosyncratic risk is consistent with recent evidence showing individuals face infrequent, large changes in income. Our estimated income shock process has a high degree of kurtosis in individual income changes. Second, individuals face uninsurable employment risk, in addition to idiosyncratic earnings risk while employed. Cyclical changes in unemployment risk imply time-varying household-level uncertainty. 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 in the Great Recession.

A third characteristic that distinguishes our analysis from most business cycle models is the overlapping-generations of households. Households start with low average levels of wealth, and accrue savings over their working lives subject to earnings and unemployment shocks. Given borrowing limits, this makes young households very sensitive to cyclical income risk. We also allow for a time-varying price of capital, which, alongside Epstein-Zin preferences, implies an environment consistent with large, countercyclical risk premia and thus a strong precautionary savings motive. Our model environment implies aggregate movements in consumption that are larger than those seen in complete markets business cycle models.

As in many quantitative incomplete markets models, in our economy, households 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. Over time they tend to accumulate wealth until retirement. Uncertain lifetimes, and a lack of annuities markets, leads them to hold substantial assets after 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.

The recent recession involved a persistent reduction in employment alongside a comparatively small and short-lived decline in TFP growth. In the context of the model, this implies a large, persistent rise in unemployment risk. As they have accumulated little savings, younger households have a higher share of income from labour, and thus experience larger lifetime effects of this persistent rise in unemployment risk on subsequent consumption. Older working households, with higher average wealth to earnings ratios, 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 lifetime consumption.

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 (2017) 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 in GDP, following large shocks to aggregate TFP. As in Krueger et al., we explore the importance of heterogeneity in an overlapping-generations environment in which households differ over their wealth and subjective discount factors. Important differences in our work include the use of Epstein-Zin preferences and a time-varying relative price of capital, which amplify the role of precautionary savings. Additionally, we study empirically consistent shocks to TFP growth.

Kim (2018) studies the effect of disaster risk in a dynamic stochastic general equilibrium model with overlapping generations. She explores how changes in aggregate disaster risk affect households' precautionary savings and thus aggregate consumption. In contrast to this work, she allows for both liquid and illiquid assets. Glover et al. (2017) study the

[1] Studies exploring the role of heterogeneous households include Krueger, Mittman and Perri (2017), Guerrieri and Lorenzoni (2015) and Glover, Heathcote, Krueger and Rios-Rull (2017). 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 (2015).

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 endogenize production and allow for capital accumulation. 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 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 a result, we follow Kim (2018) and 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, while allowing for nontrivial aggregate nonlinearities.

Section 2 of the paper describes the model, while section 3 describes 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 characterized by an initial period of work followed by retirement. In the working part of their lives, households face unemployment risk as well as earnings risk while employed. Unemployed households receive benefits and there are public pensions for retired households. Households may save or borrow subject to a limit. A third source of risk exists in uncertain longevity; households face age-specific probabilities of surviving into the next period. There are no annuities markets, and income risk is uninsurable. A representative firm employs all working households, paying each a wage proportional to their efficiency units of labour, and rents capital from an investment goods producer. Households invest their wealth with the investment firm, which produces capital subject to a technology characterized by diminishing marginal productivity of investment spending.

The aggregate state is (A_{-1}, z_i, μ) where current total factor productivity is $A^{1-\alpha} = (A_{-1}z)^{1-\alpha}$ with z a shock to TFP growth. μ is the distribution of households over the individual state space. We assume $\left\{z_i, \left\{\pi_{ij}^z\right\}_{j=1}^{N(z)}\right\}_{i=1}^{N(z)}$ is a Markov Chain with unconditional mean greater than 1. Below, we will also consider uncertainty shocks that change transition probabilities, however we abstract from such shocks in this section.

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, $\pi_{kl} \geq 0$, $k, l = 1, \ldots, N(\zeta)$, where $\sum_{l=1}^{N(\zeta)} \pi_{k,l} = 1$ for $k = 1, \ldots, N(\zeta)$. Furthermore, households of age j, $j = 1, \ldots, J-1$, survive into the next period with probability ω_j, $0 \leq \omega_j \leq 1$.

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 the duration of employment within a period. 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, here the labour force is not only determined by the aggregate exogenous state but also varies with the existing distribution of employment across households. Both earnings and unemployment shocks are independent of a household's working age, j. Households retire after $J_r - 1$ years in the labor force.

Households borrow and save using physical capital. This has a price P and pays a return D, both of which are functions of the aggregate state. The individual's state is (j, S, e, ζ) where $S \in [\underline{S}, \infty)$ is its capital and $j \in \mathbf{J} = \{1, \ldots, J\}$ describes the years since entry into the labor force. While households may borrow, their holdings of capital cannot fall below the borrowing limit, $\underline{S} \leq 0$. The budget constraint for a working household, $j = 1, \ldots, J_r - 1$, is

$$C + PS' \leq (P + D) S + (1 - \tau) W_0 \zeta l (j) e + (1 - e) B_0 \zeta \equiv M_0 (j, S, e, \zeta; A_{-1}, z_i, \mu),$$
$$C \geq 0, \; S' \geq \underline{S}.$$

Above, C is current consumption and W_0 is the real wage, which is a function of the aggregate state. When employed, households receive $(1 - \tau) W_0 \zeta l (j) e$ 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 previous income.

Labour income taxes 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,

$$C + PS' \leq (P + D) S + B \equiv M_1 (S; A_{-1}, z_i, \mu),$$
$$C \geq 0, \; S' \geq \underline{S}.$$

Social security benefits are not proportional to lifetime earnings, which captures the redistribution of the program. However they are also functions of the aggregate state, described below.

We assume that preferences, specifically subjective discount factors, vary across households. Let π_i^{β}, $i = 1, \ldots, N(\beta)$ be the fraction of households born with $\beta_i \in (0, 1)$. These subjective discount factors do not change over life.[2]

Let μ be a distribution over (i, j, S, e, ζ). Assume the distribution of households evolves according to $\mu' = \Gamma(s, \mu)$. A working household of generation $j = 1, \ldots, J_r - 1$ and type $i = 1, \ldots, N(\beta)$ solves

[2] Differences in household discount factors allow us to tractably match the distribution of wealth. The stochastic beta model of Krusell and Smith (1999) assumed random changes in stochastic discount factors to reproduce the distribution of wealth in the data. More recently, Hubmer, Krusell and Smith (2018) and Krueger Mittman and Perri (2017) assume differences in discount factors.

$$V^i\left(j,S,e_h,\zeta_k,A_{-1},z_g,\mu\right)=\max_{C,S'}\Big(\left(1-\beta_i\right)C^{1-\sigma} \tag{1}$$

$$+\beta_i\Big(\omega_j\sum_{h=1}^{N(z)}\pi_{gh}^z\sum_{i=1}^{N(e)}\pi_i^e\left(z_g\right)\sum_{l=1}^{N(\zeta)}\pi_{kl}V^i\left(j+1,S',e_i,\zeta_l,A,z_h,\mu'\right)^{1-\gamma}\Big)^{\frac{1-\sigma}{1-\gamma}}\Big)^{\frac{1}{1-\sigma}}$$

subject to

$$C+PS'\leq M_0\left(j,S,e,\zeta;A_{-1},z_i,\mu\right),$$

$$\mu'=\Gamma(s,\mu)$$

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.

We describe retired households that entered the labor force $j=J_r,\ldots,J-1$ periods ago. Such households, of type $i=1,\ldots,N\left(\beta\right)$, solve the following borrowing and savings problem.

$$V^i\left(j,S,e_h,\zeta_k,A_{-1},z_g,\mu\right)=\max_{C,S'}\Big(\left(1-\beta_i\right)C^{1-\sigma} \tag{2}$$

$$+\beta_i\Big(\omega_j\sum_{h=1}^{N(z)}\pi_{gh}^z\sum_{i=1}^{N(e)}\pi_i^e\left(z_g\right)\sum_{l=1}^{N(\zeta)}\pi_{kl}V^i\left(j+1,S',e_i,\zeta_l,A,z_h,\mu'\right)^{1-\gamma}\Big)^{\frac{1-\sigma}{1-\gamma}}\Big)^{\frac{1}{1-\sigma}}$$

subject to

$$C+PS'\leq M_1\left(S;A_{-1},z_i,\mu\right).$$

$$\mu'=\Gamma(s,\mu)$$

Retired households only face aggregate risk; they have no idiosyncratic earnings or employment risk. Regardless, we include e and ζ in their value functions for consistency with (1) when $j=J_r-1$.

Households that reach J years of age since labor-force entry, the last possible age, simply consume their cash on hand.

$$V^i\left(J,S,e_h,\zeta_k,A_{-1},z_g,\mu\right)=\Big(\left(1-\beta_i\right)\left(M_1\left(J,S\right)\right)^{1-\sigma}\Big)^{\frac{1}{1-\sigma}}. \tag{3}$$

2.2 Production, investment 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 $\mu\left(i,j,S,e,\zeta\right)$ describe the measure of households of working age j with wealth S, employment status e, and labour productivity ζ. The aggregate labour supply is the sum of working-age households, over labour productivity and employment,

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

The aggregate capital stock is

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

A competitive final goods firm produces using the technology described by

$$Y = \left(A_{-1} z_g\right)^{1-\alpha} K^\alpha L^{1-\alpha}$$

where $0 < \alpha < 1$. The firm solves the following problem,

$$\max_{K,L} \Big(\left(A_{-1} z_g\right)^{1-\alpha} K^\alpha L^{1-\alpha} - r_0^k \left(A_{-1}, z_g, \mu\right) K - W_0 \left(A_{-1}, z_g, \mu\right) L \Big).$$

where $r_0^k \left(A_{-1}, z_g, \mu\right)$ is the rental rate for capital. In equilibrium, capital and efficiency units of labour satisfy

$$r_0^k \left(A_{-1}, z_g, \mu\right) = \alpha \left(A_{-1} z_g\right)^{1-\alpha} K^{\alpha-1} L^{1-\alpha}, \tag{4}$$

$$W_0 \left(A_{-1}, z_g, \mu\right) = (1-\alpha) \left(A_{-1} z_g\right)^{1-\alpha} K^\alpha L^{-\alpha}. \tag{5}$$

Capital is rented from a firm that invests household wealth and produces capital. This investment goods firms problem solves

$$\begin{aligned} F^I \left(K, A_{-1}, z_g, \mu\right) = \max_{K'} \Bigg(& \left(r_0^k + 1 - \delta\right) K - \Big(P + D\Big) K \\ & + P K' - \psi\left(K', K\right) \\ & + \sum_{h=1}^{N(z)} \pi_{gh}^z q\left(z_h, A_{-1}, z_g, \mu\right) F^I \left(K', A, z_h, \mu'\right) \Bigg). \end{aligned} \tag{6}$$

The firm's minimum cost of producing K', given an existing capital stock of K, is given by the function $\psi\left(K', K\right)$. It pays households a total return of $(P + D)$, which it takes as given, on the capital they deposit with it, and sells them new capital K' at the price P. In equilibrium, the competitive investment firms demand for K and production of K' satisfy

$$P\left(A_{-1}, z_g, \mu\right) = \psi_1\left(K', K\right)$$
$$D\left(A_{-1}, z_g, \mu\right) = r_0^k\left(A_{-1}, z_g, \mu\right) + 1 - \delta - P\left(A_{-1}, z_g, \mu\right) - \psi_2\left(K', K\right).$$

Assuming ψ is a constant returns to scale function, this implies an equilibrium firm value of 0 in every state, which allows us to avoid specifying the stochastic discount factor, $q\left(z_h, A_{-1}, z_g, \mu\right)$.

As already mentioned, tax revenues are used to fund social security, unemployment benefits and government spending. Define the total population of age j as

$$\mu_j = \sum_{i=1}^{N(\beta)} \pi_i^\beta \sum_{j=1}^{J_r} \int_{[\underline{S},\infty)\times\left\{e_1,\ldots,e_{N(e)}\right\}\times\left\{\zeta_1,\ldots,\zeta_{N(\zeta)}\right\}} \mu\left(i, j, d\left[S \times e \times \zeta\right]\right),$$

the government budget constraint is

$$b \sum_{k=1}^{N(\zeta)} \pi_k^0 \zeta_k \sum_{j=J_r}^{J} \mu_j^a + \left(1 - \pi_1^e\right) b_0 \sum_{k=1}^{N(\zeta)} \pi_k^0 \zeta_k \sum_{j=1}^{J_r-1} \mu_j^a + G = \tau w L + S^a \tag{7}$$

where G is government spending and S^a is the sum of capital saved by households that do not survive,

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

2.3 The distribution of households

We assume that initial wealth of newly working households is 0, and that their distribution over employment and labour productivity 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 $\{\pi_{kl}\}_{k,l=1}^{N(\zeta)}$. We assume that a constant number of new households, $\sum_{i=1}^{N(\beta)} \pi_i^\beta \pi_h \int_{[\underline{S},\infty)\times\left\{e_1,\ldots,e_{N(e)}\right\}\times\left\{\zeta_1,\ldots,\zeta_{N(\zeta)}\right\}} \mu\left(i, 1, S, e, \zeta\right)$, enters the labour force each period. Since survival probabilities are independent of the aggregate state, this implies a constant mass of households in the economy.

Given $\mu\left(i, 1, S, e, \zeta\right)$, the distribution of newly working households over employment, labour productivity and wealth, the evolution of households over age is given by

$$\mu\left(i, j+1, B, e_h, \zeta_k\right) = \omega_j \pi_h \sum_{n=1}^{N(\zeta)} \pi_{nk} \int_{\{(S,e,\zeta)|g(j,S,e,\zeta,s)\in B\}} \mu\left(i, j, d\left[S \times e \times \zeta\right]\right) \tag{8}$$

$$j = 1, \ldots, J-1,\ i = 1, \ldots, N(\beta) \text{ and } h = 1, \ldots, N(e).$$

where $g(i, j, S, e, \zeta, A_{-1}, z_i, \mu)$ is the savings decision rule for a household of type i, age j, beginning of period wealth S, current employment status e, and labour productivity ζ, given an aggregate state (A_{-1}, z_i, μ).[3] The distribution $\mu(i, j, S, e, \zeta)$ is the endogenous component of the aggregate state. We study its evolution in recursive general equilibrium with aggregate shocks.

2.4 Recursive Competitive Equilibrium

A *recursive competitive equilibrium* is a set of functions,

$$\{\{V_i\}_{i=1}^{N(\beta)}, F_I, g, K', K, L, r_o^k, W_o, P, D, B, B_0\}$$

that solve the problems of households, goods-producing firms, and capital-producing firms. These functions clear the markets for assets, labor, and output, by satisfying the following conditions:

1. V_i solves (1)-(3) for every $i \in \{1 \ldots, N(\beta)\}$, with g being the associated policy function
2. F_I solves (6), with K' being the associated policy function
3. K and L maximise profits for the final goods firm
4. Γ satisfies (8)
5. Markets Clear

3 Parameters

We set the length of a period to one year, and assume that households enter the labor force at age 25. Given this, we set $J_r = 40$ and $J = 70$, assuming that retirement is at age 65 and no households live longer than 95. Given this life-cycle, we assume that $\omega_j = 1$ for

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

$j = 1, \ldots, J_r - 1$, then falls linearly until J it is 0. Thus, in this preliminary version of the paper, rather than calibrate to data on mortality risk, we simply assume

$$\omega_j = \frac{J - j}{J + 1 - J_r} \text{ for } j = J_r, \ldots, J.$$

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 11. These values go a long way in generating reasonable equity premia and a risk-free real interest rate in a complete markets model with long-run TFP risk.

In the example we study in the remainder of the paper, we assume $N(\beta) = 3$ and choose a support for β_i, (0.9, 0.98, 0.995). The percentage of households born with $\beta_i = 0.9$ and 0.98 is 0.85 and 0.10, respectively. The heterogeneity in β increase the concentration of wealth in our example to roughly 40 percent of that seen in the data reported in Table 3 below. If we abstract from household-level differences in β, the share of wealth held by the top one percent of households falls to less than 15 percent of that in the data.[4] Our distribution of β, given the income process described below, implies a 4.3 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 $B = 0.45w$.

3.1 Earnings Shocks Estimation

We estimate a leptokurtic income shocks process following Kaplan, Moll and Violante (2016). Our income shocks reproduce 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

[4] We are currently working to better match the distribution of wealth.

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). However, our parameters are estimated using simulated methods of moments to minimize, for our annual calibration, the sum of squared residuals with the eight empirical moments above.

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

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

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, we rely on 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, 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 z, 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). Our TFP growth process is assumed to be log-normal. We assume that the underlying normal process is AR(1) with a persistence of 0.909 and a standard deviation of innovations of 0.0109. In our model economy, with exogenous variation in total hours worked arising from countercyclical unemployment risk, this will imply a percent standard deviation of detrended output comparable to the data.

3.2 Solution

We discretize individual earnings and solve decision rules The leptokurtic income shocks process, estimated to match the moments from the SSA data, is discretized using a 25 million period simulation. We choose 15 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.

Household value functions are solved using a generalization of the Backwards Induction algorithm of Reiter (2010) described by Kim (2018), 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 abstracts 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 (z,K), given a proxy distribution $\mu(j, e, \zeta, S; z, K)$, we conjecture a value for $\widehat{K}'$, the future approximate aggregate state.[5] Next, we determine the actual end-of-period aggregate capital stock by solving household value functions at each age, savings level, 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 $v^{n+1}(J, e, \zeta, S, z, K)$ using (3) where n indexes the iteration of value functions, across ages, with respect to the aggregate state (z, K). Using $v^{n}(j+1, e', \zeta', S', z', K')$, $j = 1,..., J-1$, we solve decision rules using the endogenous grid method of Carroll (2005). Alongside $v^{n+1}(j, e, \zeta, a, z', k')$ this provides decision rules $S' = g^{n+1}(j, e, \zeta, S, z, K)$ 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 minimization 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 (2018)

[5] We solve a stationary model by deflating household and firm problems by $A_{=1}$, this eliminates the level of last period's total factor productivity from the state space.

provides a feasible reduction, aggregating over (j, e, ζ) using weights $\omega_{j,e,\zeta}(S)$, 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 1.88 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 the leptokurtic income shock process. Over time they borrow or save, and their wealth grows. Over their life-cycles, on average households accumulate wealth. Thus younger households tend to be poorer than older working households and 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	0.12	0.36	0.51	0.91	1.0	0.64	0.09

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 13 percentage points of the SCF data in 2004. Nonetheless, in the absence of a stronger motive to accumulate wealth over and beyond consumption smoothing and retirement, there is less skewness in its distribution than in the data. Thus the wealthiest 1 percent of households hold 12 percent of the total wealth of the economy, while, in the data, the corresponding number is 29 percent. However, the overlapping generations economy, with an estimated income shock process, is able to generate an unusually large level of inequality. Half of the inequality in the model arises from income and survival risk, the

Figure 1: The Distribution of Wealth

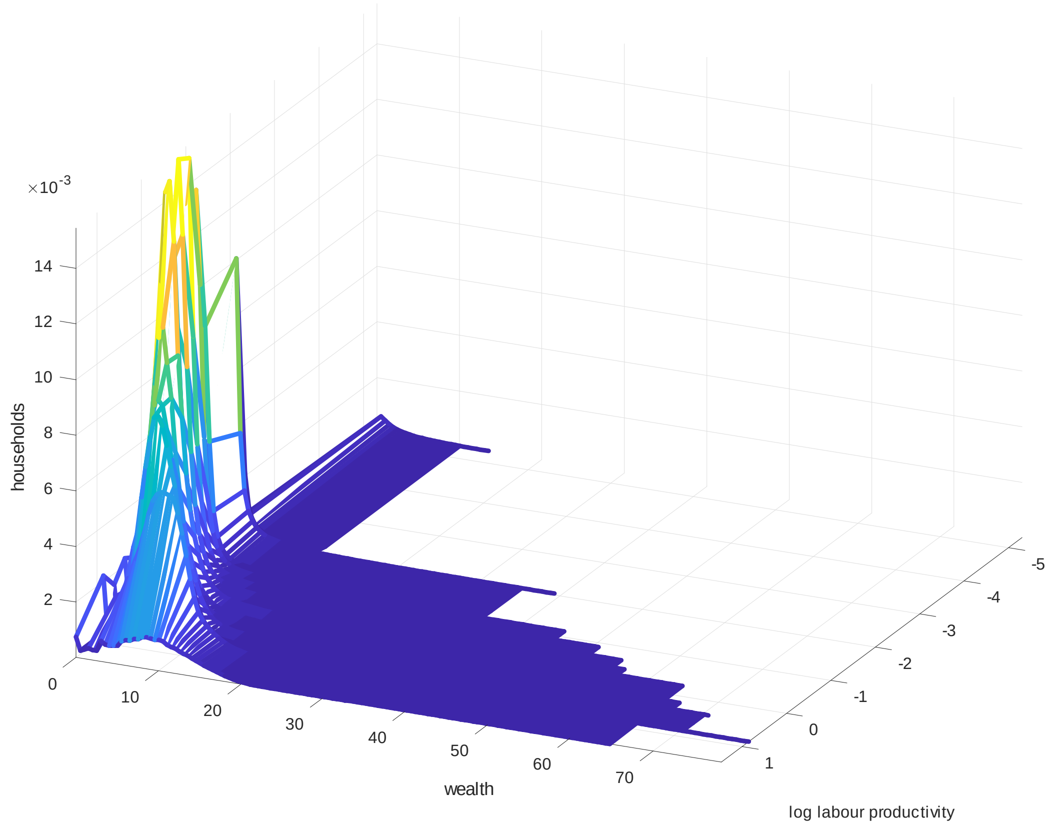

rest is driven by life-cycle wealth accumulation. When we consider an alternative economy with no earnings risk, while retaining employment risk, and deterministic lifetimes, the Gini coefficient falls to 0.34 and the richest one percent of households own 3.25 percent. of the capital stock. In an infinite horizon model there would be almost no inequality as unemployment risk is idiosyncratic and short-lived. In such models, all differences in wealth arise from income shocks. Here, life-cycle savings drive inequality.

Figure 1 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 and higher income households tend to be wealthier given the persistence in the estimated income shock process.

Figure 2 shows the average levels of income, consumption and wealth in a generation over time. Its members retire after 45 years of work. As mentioned in the beginning of section 3, the social security replacement rate used in the model is 0.40 of average earnings, which is high. As this is independent of individual earnings, differences in individual income

Figure 2: A Cohort of Households

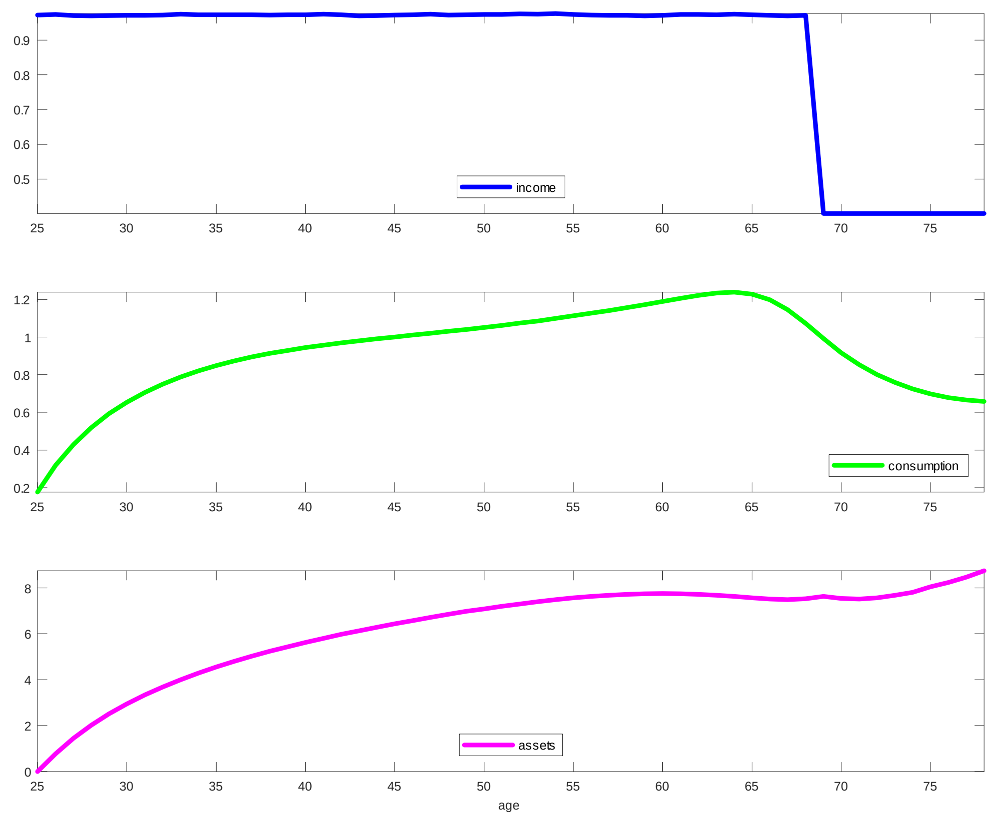

disappear. This income, common to all households, is shown in the top panel starting with age 70.

The equilibrium real interest rate in the steady state is 4.3 percent, while households' subjective discount factors vary between 0.9 and 0.995. Income risk and life-cycle savings drive the accumulation of wealth seen in the bottom panel. On average, households save until retirement. Thereafter, uncertain lifetimes maintain average wealth in the cohort, which actually shows a slight rise ten years after retirement. Survival risk, after retirement, offsets the incentive to dissave given finite lifetimes and most households leave accidental bequests that are collected by the government. While working, the average rise in income and wealth implies a gradual increase in consumption. Consumption falls gradually after retirement, as households increase their savings rates to maintain their wealth. This fall in consumption later is the result of an increase in precautionary savings which, after retirement, is used to smooth consumption against an uncertain lifespan.

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 roles of income and wealth

Figure 3: Cohort standard deviations over age

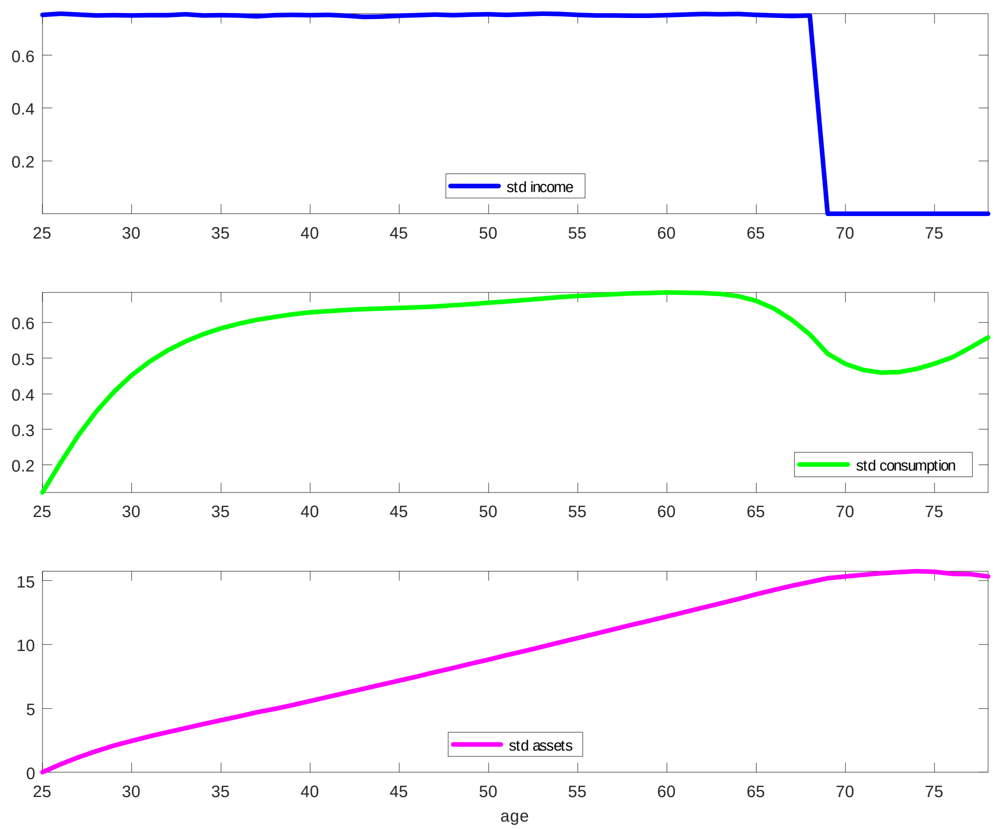

heterogeneity, and age, in determining the equilibrium response to aggregate shocks.

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 growth 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	MPK	w
std	1.87	2.00	1.12	3.16	1.68	1.92	0.22	1.17
correlation	0.73	1.00	0.75	0.97	0.79	0.04	0.93	0.55

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 departures from the latter. First, incomplete markets with a strong precautionary savings motive, implied by the high risk aversion characterizing our Epstein-Zin preferences, affects households' willingness to intertemporally substitute consumption. We see this by comparing our economy to one with expected utility where the elasticity of intertemporal substitution is 0.5 ($\sigma = 2$). In this alternative setting, we maintain all parameter values except the support of the distribution of β_i across households, which we uniformly raise by 10.5 percent to match the same average marginal product of capital, of 4.3 percent, as in our benchmark economy with Epstein-Zin utility. The economy with constant relative risk aversion has a consumption percent standard deviation of 2.11, almost twice that of the economy studied here. The economy without high risk aversion has less precautionary savings, and a Gini coefficient of 57 percent, 7 percentage points less than our benchmark.

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 endogenous labor supply.

[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 marginal product of capital (MPK) and the real wage (w). The first row reports percent standard deviations while the second reports the contemporaneous correlation of each series with output.

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. 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 0.55, 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 GDP in the United States fall by 9 percent through the trough in 2009Q2, relative to its long-run growth trend. Consumption, detrended using its long-run growth rate, fell 6.5 percent while private investment dropped by 29 percent and total hours worked fell 9.5 percent.[7] Surprisingly, measured TFP fell by only 3.3 percent over the same period. We explore the impact, on the model economy, of empirically consistent shocks to TFP growth and employment.

Our exercise involves a persistent shock to TFP growth rates, of a magnitude seen in the Great Recession. This drives a rise in employment risk and a fall in total hours worked. The unemployment rate, which includes partial unemployment for any part of a year, initially reaches almost 12 percent. We assume that TFP growth has a persistence of 0.91, this is similar to the persistence of the Solow residual measured by Khan and Thomas (2013). As next period's employment risk is determined by current total factor productivity growth, there is a one period lag in employment. Figure 4A shows the growth shock and the effect of the rise in employment risk on the labour input. Notice that total hours worked fall by roughly 7 percent, which is somewhat less than the decline reported in Figure 4 above.

The mechanics of a shock to TFP growth are distinct from the familiar response to a persistent shock when the process is stationary. As the growth shock is persistent, and involves a period of negative growth rates as seen above, the level of TFP has a hump-shaped response (Figure 4B). This reduces the substitution effect that implies, when shocks are stationary, a hump-shaped response in aggregate consumption. A negative shock to TFP growth implies a period of falling total factor productivity and thus a return to savings that falls over time. Lower future returns reduce the substitution effect that, in a model with monotone responses in aggregate TFP, leads to a gradual, hump-shaped response in consumption. As seen in Figure 5, the percent response in consumption, following a growth

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

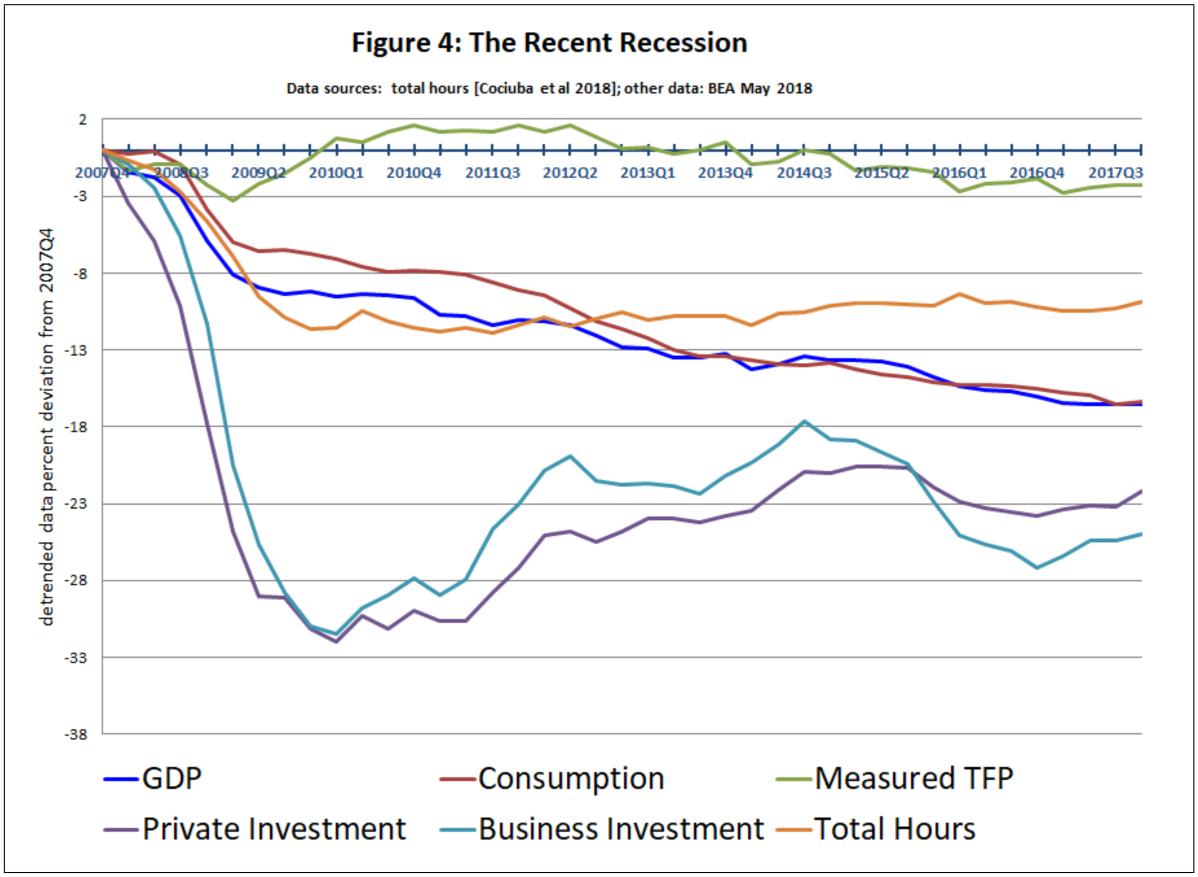
Figure 4: The Recent Recession
Data sources: total hours [Cociuba et al 2018]; other data: BEA May 2018
detrended data percent deviation from 2007Q4
2
-3
-8
-13
-18
-23
-28
-33
-38
2007Q4
2008Q3
2009Q2
2010Q1
2010Q4
2011Q3
2012Q2
2013Q1
2013Q4
2014Q3
2015Q2
2016Q1
2016Q4
2017Q3
GDP
Consumption
Measured TFP
Private Investment
Business Investment
Total Hours

Figure 4A: Growth and Employment Shocks

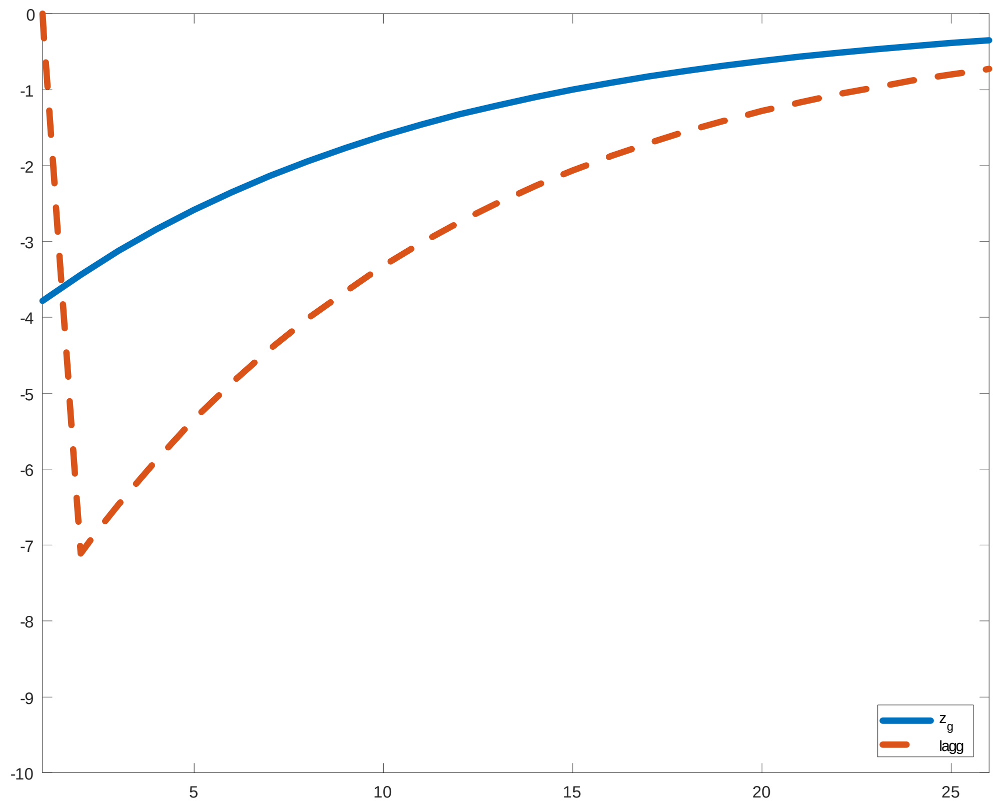

Figure 4B: Change in Total Factor Productivity

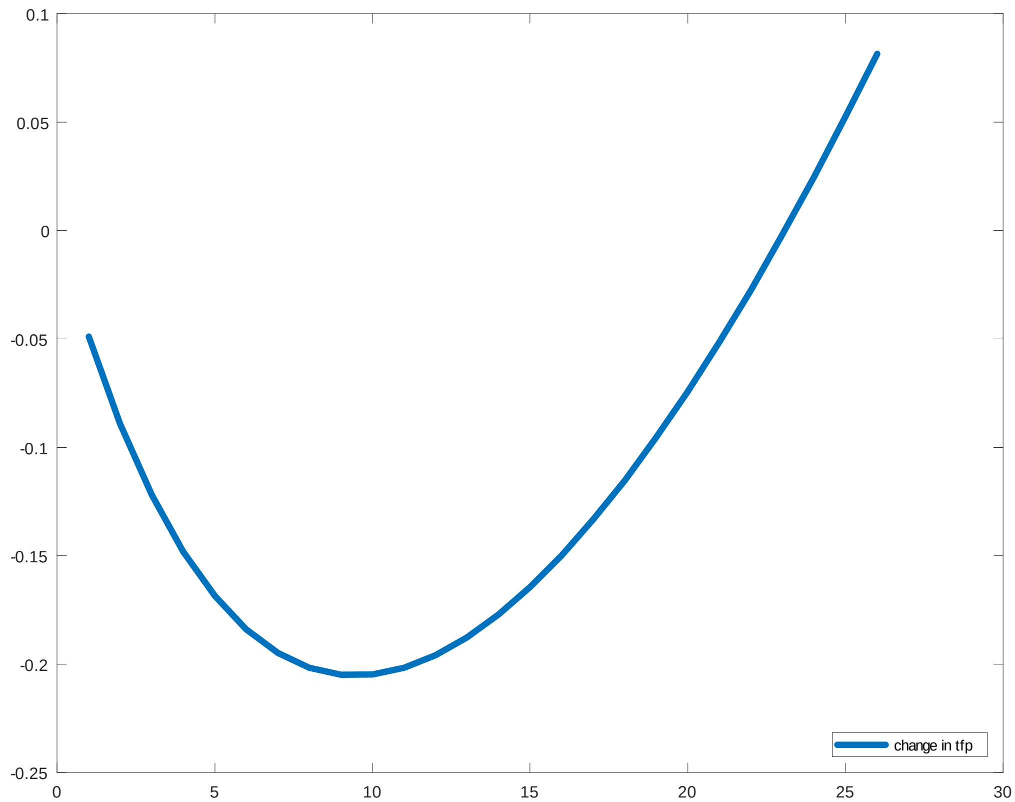

shock, is monotone and eventually of a magnitude similar to GDP. The change in the rate of fall, in consumption, investment and GDP, one period after the shock is the result of the delayed fall in employment. Employment lags because the unemployment risk varies with lagged TFP.

As there is a one-period lag in employment, we compare the change in aggregate quantities in our annual model after two periods to the peak-to-trough changes seen in the data. In the model economy, the persistent shock to TFP growth leads to a 9.2 percent fall in GDP, against trend, and a 4.4 percent fall in consumption. Aggregate investment falls 12.9 percent and total hours worked decline 7.1 percent. As the shock is persistent, the growth recession continues past two periods.

Thus the model is able to explain all of the fall in GDP and 68 percent of the drop in consumption. The results for investment are less strong; we generate 45 percent of the peak-to-trough change. 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.

Figure 5: GDP, Consumption and Investment

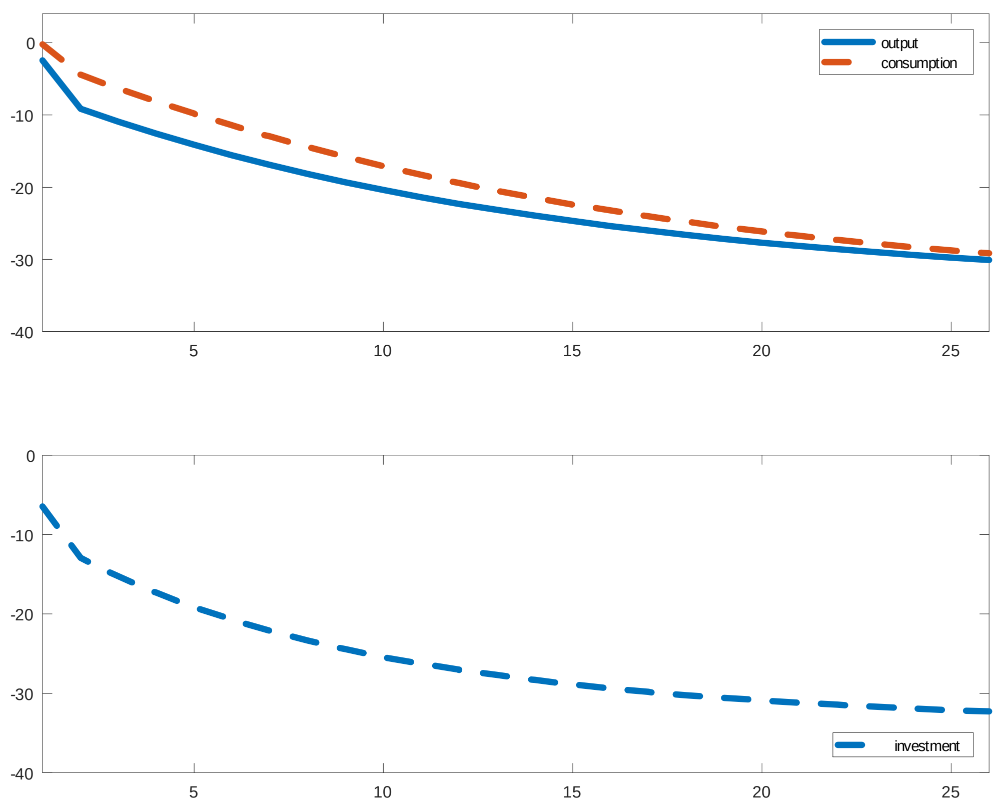

Importantly, if we are to explain the fall in employment, which is difficult to explain from a frictionless representative agent model with variable labour supply (see Khan and Thomas 2013), we have to introduce a rise in unemployment risk. This implies a large welfare cost of the recession on younger households. Moreover, the resulting fall in total hours amplifies and propagates the fall in TFP growth.

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.

Our next step is to introduce time-varying volatility of TFP growth. This will allow us to address the slow recovery in aggregate quantities alongside the rebound in TFP growth seen in Figure 4. We are also recalculating the welfare costs of the recession over cohorts. Initial results indicate larger declines in welfare for young workers and retired households, with smaller costs of the recession on middle-aged workers. The latter are better able to insure their consumption relative to young workers who have less wealth. As their horizon is longer than that of retirees, they also benefit a recovery in the price of capital which determines the value of their assets. Retired households see a large fall in their consumption, and thus higher welfare costs, as a result of the price of capital.

References

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

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

[3] Cociuba, Simona E., Edward C. Prescott, and Alexander Ueberfeldt (2018) “US hours at work.” *Economics Letters* 169, pages 87-90.

[4] 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.

[5] Glover, Andrew, Jonathan Heathcote, Dirk Krueger and Jose Victor Rios-Rull (2017) ‘Intergenerational Redistribution in the Great Recession,’ working paper.

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

[7] 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.

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

[9] 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.

[10] Hubmer, Joachim, Per Krusell and Anthoy A. Smith (2018) ‘A Comprehensive Theory of the U.S. Wealth Distribution.’ working paper.

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

[12] Kim, Heejeong (2018) ‘Inequality, Portfolio Choice and the Business Cycle’ working paper.

[13] 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.

[14] 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.

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

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

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