The Persistence of Recessions with Incomplete Markets and Time-Varying Income Risk

Aubhik Khan

Ohio State University

Ben Lidofsky[1]

Ohio State University

15 February 2022

ABSTRACT

We study the propagation of recessions in overlapping generations economies wherein households, with uncertain lifetimes and uninsurable earnings risk, face cyclical employment risk. Business cycles are driven by persistent shocks to TFP growth and household-level employment. Increases in employment risk cause fluctuations in both the unemployment rate and in labor force participation. In this setting, we introduce elements commonly used to deliver a strong and countercyclical precautionary savings motive. Specifically, households have non-separable utility characterised by high levels of risk aversion, and a diminishing marginal productivity of investment leads to a time-varying price of capital.

We find that changes in precautionary savings, following aggregate shocks, have important implications for aggregate consumption. Persistent negative shocks to TFP growth, associated with increases in risk to employment, drive large declines in consumption. This helps explain the large fall in consumption observed over the Great Recession. An empirically consistent, moderate shock to TFP growth rates implies a large and persistent fall, against trend, in aggregate consumption. Moreover, an estimated rise in households' risk of long-term non-employment reduces labor force participation and reconciles the swift recovery in TFP growth rates with a protracted decline in consumption and output.

[1]Aubhik Khan: mail@aubhik-khan.net; Ben Lidofsky: benjamin.lidofsky@gmail.com. We thank seminar participants at Iowa State University, the National University of Singapore, the University of Georgia and Georgetown University as well as conference participants at the Midwest Macro Meetings. We thank Mark Huggett, Toshi Mukoyama and Serene Tan for helpful comments.

1 Introduction

Quantitative business cycle theory has struggled to explain the volatility of consumption. This presents a particular challenge when considering the Great Recession, which featured an unusually long and large decline in aggregate consumption in the United States. We hypothesize that elevated idiosyncratic risk and precautionary savings motives across individual households are primary drivers of the aggregate trend. We study this using an overlapping generations model in which individuals face uninsurable risk to their earnings and labor force status. The economy as a whole is also exposed to uncertainty, with total factor productivity subject to trend-stationary growth shocks and labor market varying risk across the business cycle. Consequently, our households are simultaneously faced with several sources of risk that they need to save against, a combination that has not been examined before. These forces lead to strong cyclical variation in precautionary savings behaviour, which in turn amplifies the volatility in aggregate consumption.

A notable example of a large movement in aggregate consumption is the Great Recession, in which a mild, brief fall in measured TFP coincided with unusually large, long declines in aggregate output and consumption. We find that trend shocks to TFP cannot explain the weak recovery in consumption on their own. However, the Great Recession was also a time of heightened labor market risk. When we confront our households with empirically consistent rises in employment risk and shocks to TFP, a substantial rise in precautionary savings generates an initial fall in aggregate consumption much closer to what we see in the data. However, this rise in precautionary savings is strong enough to create a counterfactually mild fall in aggregate investment, which in turn dampens the longer-run decline in aggregate output. We show that a shock to investment productivity, combined with heightened labor market risk, produces both large declines in aggregate investment and a prolonged fall in aggregate consumption, just as we saw following the Great Recession.

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. 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 risk of unemployment and labor-force exit imply time-varying household-level hours worked. Examining the effects of a large recession, we account for the changes in labor force composition during the Great Recession. Using data from the Survey of Income and Program Participation, we identify a reduced probability of people transitioning into the labor force relative to previous recessions. Accounting for this, our labor market dynamics generate a sustained fall in labor force participation that, combined with a rise in unemployment, implies both heightened risk for individual workers and a weaker labor supply in the aggregate.

As in many quantitative incomplete markets models, in our economy households borrow and save to smooth their consumption against fluctuations in their income, and are subject to a borrowing limit. A third characteristic that distinguishes our analysis from most business cycle models is the overlapping generations of households. In our setting, finite lifetimes imply that cyclical variation in income leads to a permanent change in lifetime earnings. Consequently, households respond to any fall in income with a permanent change in their lifetime consumption. This amplifies the effect of business cycles on aggregate consumption.

Overlapping generations models also introduce life-cycle sources of inequality beyond what arises from income shocks alone. Our households age discretely, working for a known set of periods before transitioning into a retirement of uncertain duration, owing to mortality risk. In our model, young households are wealth-poor and, on average, choose 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 a standard incomplete markets model with infinitely lived households. Furthermore, as households start with low average levels of wealth, and accrue savings over their working lives subject to earnings and unemployment shocks, young households are 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.

The observed movements in aggregate quantities over the Great Recession present an empirical anomaly over and beyond the size of the recession itself. These have proven difficult to explain using canonical representative agent business cycle models driven by persistent shocks to the level of exogenous TFP. Khan and Thomas (2013) study credit shocks in a model with a time-varying distribution of firms. Their model with production heterogeneity can explain the large fall in GDP and investment alongside the relatively modest reduction in TFP. However, their complete markets model is unable to address the slow recovery observed in the data and the large fall in consumption.

Given the limited success of complete markets models to explain the large fall in consumption, we turn in this work to a model with incomplete markets and a time-varying distribution of households. As discussed above, we find that growth shocks, when combined with increases in employment risk, can reconcile the small fall in TFP with large declines in aggregate consumption and GDP, and slow recoveries. Importantly, in our setting, the shock to TFP growth is short-lived. The slow recovery is primarily driven by a persistent rise in the risk of non-employment, which reproduces the long-term decline in labor force participation seen in the data. Furthermore, this rise in idiosyncratic risk, in our setting where high levels of risk aversion characterise household preferences, leads to a large fall in aggregate 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 (forthcoming) 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. (2020) 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 aggre-

[1] Studies exploring the role of heterogeneous households include Krueger, Mittman and Perri (2017), Guerrieri and Lorenzoni (2017) and Glover, Heathcote, Krueger and Rios-Rull (2020). 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).

gate 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. While, in our model, the incidence of a recession varies over age as well as income and wealth, it does so differently from the findings of Glover et al. (2020). First, our production economy is subject to the large, persistent rises in non-employment risk observed in the data. As they have accumulated little savings, younger households have a higher share of income from labor, and thus experience larger lifetime effects of this persistent rise in non-employment risk on subsequent consumption. Older working households, with higher average wealth to earnings ratios, are better able to hedge against joblessness. Nonetheless, finite working lives imply larger welfare costs of persistent increases in non-employment. A persistent fall in average income across households leads to a permanent drop in lifetime consumption. In other related work, Guerrieri and Lorenzoni (2017) 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 shocks to employment risk 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 (forthcoming) 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 non-employment risk as well as earnings risk while employed. Households receive benefits when they are not employed 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 labor, 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. Households' transition probabilities over labor-force states is also subject to aggregate risk. We assume that such risk is perfectly correlated with shocks to z, and hence this is not a standalone element of the aggregate state.

2.1 Households

Households that have been working j periods have experience $l(j)$. Each household has a persistent idiosyncratic labor 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 labor market risk that is uncorrelated with their labor productivity. Let $e \in \left\{e_1, \ldots, e_{N(e)}\right\}$ describe the set of feasible employment durations within a period, where $1 \geq e_1 > e_2 > \cdots > e_{N(e)} \geq 0$. We assume that labor market risk follows a first-order Markov process, with transition probabilities $\pi_{i,j}^e(s) \geq 0,\ i,\ j \in \{1, \ldots, N(e)\}$ with $\sum_{j=1}^{N(e)} \pi_{i,j}^e(s) = 1\ \forall\ i$. The exogenous component of the aggregate state is s, and idiosyncratic labor market risk in the economy may change with aggregate shocks as in Krusell and Smith (1999). However, in contrast to that paper, here the labor 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 $J_r - 1$ periods 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 periods 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

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

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 labor income tax rate. Non-employed households receive benefits, $B_0 \zeta$ that are proportional to their labor productivity. The persistence of labor productivity implies that non-employment benefits vary with previous income.

Labor income taxes finance government spending, non-employment benefits and social security payments to retired households, B. These households, in cohorts $j = J_r, \ldots, J$, have the following budget constraint,

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

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]

[2]Differences in household discount factors allow us to tractably match the distribution of wealth. In an infinite-horizon setting, 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.

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

$$V^i(j, S, e_k, \zeta_l, A_{-1}, z_g, \mu) = \max_{C,S'} \Big((1-\beta_i) C^{1-\sigma} \tag{1}$$
$$+\beta_i \Big(\omega_j \sum_{h=1}^{N(z)} \pi^z_{g,h} \sum_{m=1}^{N(e)} \pi^e_{k,m}(z_g) \sum_{n=1}^{N(\zeta)} \pi_{l,n} V^i \left(j+1, S', e_m, \zeta_n, A, z_h, \mu'\right)^{1-\gamma} \Big)^{\frac{1-\sigma}{1-\gamma}} \Big)^{\frac{1}{1-\sigma}}$$
subject to
$$C + PS' \leq M_0(j, S, e, \zeta; A_{-1}, z_i, \mu),$$
$$C \geq 0, \; S' \geq \underline{S}$$
$$\mu' = \Gamma(s, \mu)$$

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 $i = 1, \ldots, N(\beta)$, solve the following borrowing and savings problem.

$$V^i(j, S, e_k, \zeta_l, A_{-1}, z_g, \mu) = \max_{C,S'} \Big((1-\beta_i) C^{1-\sigma} \tag{2}$$
$$+\beta_i \Big(\omega_j \sum_{h=1}^{N(z)} \pi^z_{g,h} \sum_{m=1}^{N(e)} \pi^e_{k,m}(z_g) \sum_{n=1}^{N(\zeta)} \pi_{l,n} V^i \left(j+1, S', e_m, \zeta_n, A, z_h, \mu'\right)^{1-\gamma} \Big)^{\frac{1-\sigma}{1-\gamma}} \Big)^{\frac{1}{1-\sigma}}$$
subject to
$$C + PS' \leq M_1(S; A_{-1}, z_g, \mu).$$
$$C \geq 0, \; S' \geq \underline{S}$$
$$\mu' = \Gamma(s, \mu)$$

Retired households only face aggregate risk; as they no longer work, they are not subject to 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 periods of age since labor-force entry, the last possible age, simply consume their cash on hand.

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

2.2 Production, investment and government

A representative firm uses capital and labor to produce output. We begin by describing the aggregate supply of efficiency units of labor. 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 labor productivity ζ. The aggregate labor supply is the sum of working-age households, over labor 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,

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

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

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

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

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

$$F^I\left(K, A_{-1}, z_g, \mu\right) = \max_{K'}\Bigg(\left(r_0^k + 1 - \delta\right)K - \Big(P + D\Big)K \tag{6}$$

$$+ PK' - \psi\left(K', K\right)$$

$$+ \sum_{h=1}^{N(z)} \pi_{gh}^z q\left(z_h, A_{-1}, z_g, \mu\right) F^I\left(K', A, z_h, \mu'\right)\Bigg).$$

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

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

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 labor 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 labor productivity, determined by $\{\pi_{kl}\}_{k,l=1}^{N(\zeta)}$. We assume that a constant number of new households, $\sum_{i=1}^{N(\beta)} \pi_i^\beta \pi_h \int_{[\underline{S},\infty)\times\{e_1,\ldots,e_{N(e)}\}\times\{\zeta_1,\ldots,\zeta_{N(\zeta)}\}} \mu(i,1,S,e,\zeta)$, enters the labor force each period. Since survival probabilities are independent of the aggregate state, this implies a constant mass of households in the economy.

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

$$\mu(i,j+1,B,e_h,\zeta_k) = \omega_j \pi_h \sum_{n=1}^{N(\zeta)} \pi_{nk} \int_{\{(S,e,\zeta)|g(j,S,e,\zeta,s)\in B\}} \mu(i,j,d[S\times e\times\zeta]) \tag{8}$$
$$j=1,\ldots,J-1,\; i=1,\ldots,N(\beta) \text{ and } h=1,\ldots,N(e).$$

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 labor 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

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

4. Γ satisfies (8)

5. Markets clear

3 Parameters

We set the length of a period to a quarter of a year, and assume that households begin working in their 25th year. Given this, we set $J_r = 180$ and $J = 280$, assuming that retirement is at age 65 and no households live longer than 95 years. We assume that working-age households face no mortality risk. For retirement-age households, we use data from the Centers for Disease Control and Prevention to discipline our one-period-ahead survival probabilities.

In developing a framework that allows for realistic risk premia, we allow for a fairly high elasticity of intertemporal substitution and set $\sigma = \frac{2}{3}$. We set $\gamma = 11$ to generate a coefficient of relative risk aversion of 16.5. 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 $N(\beta) = 2$ and choose a support for β_i, $\{0.91, 0.999999\}$ annualized. 90% of households are born with the lower discount rate. The heterogeneity in β increases the concentration of wealth substantially. Relative to a common-β environment, β heterogeneity increases the Gini coefficient by a third and roughly doubles the concentration of wealth among the top one percent of households. Our distribution of β, given the income process described below, implies a 5.5 percent average real interest rate at an annual frequency.

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 labor tax rate is $\tau = 0.3$. Unemployment benefits are 43.5 percent of earnings, $B_0 = 0.435W_0$. The social security replacement rate is $B = 0.40W_0$.

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. (2021).

Table 1: Individual Earnings Process

moment	data	continuous	discrete
std 1 year change	0.51	0.59	0.69
std 5 year change	0.78	0.68	0.76
skewness 1 year change	-1.07	-0.04	-0.29
skewness 5 year change	-1.25	0.01	0.00
kurtosis 1 year change	14.93	13.90	13.95
kurtosis 5 year change	9.51	9.15	9.29
fraction 1 year change $< 5\%$	0.306	0.300	0.294
fraction 1 year change $< 10\%$	0.488	0.633	0.463
fraction 1 year change $< 20\%$	0.665	0.823	0.676

We assume an income process that is similar to that estimated by Kaplan et al. (2017), however our parameters are estimated using simulated methods of moments to minimise 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 quarterly model, we estimate the 6 parameters using a simulation lasting 4 million periods. The resulting match to the data are shown above in Table 1 and the estimated parameters values are reported in Table 2.

3.2 Estimation of Labor Force Process

We assume three labor force states: employed (E), unemployed (U), and not in the labor force (N). Labor force transitions follow a first-order Markov process, and vary with the

Table 2: Income Process Parameter Estimates

parameter	ρ_1	ρ_2	σ_1	σ_2	λ_1	λ_2
value	-0.16	0.96	1.74	0.40	0.09	0.07

exogenous aggregate state. In particular, we estimate different transition probabilities for each aggregate TFP state. Additionally, there is a separate transition matrix when a high uncertainty shock coincides with a negative TFP shock. We rely on this last transition matrix to capture the unusual labor-force trends observed over the Great Recession.

3.2.1 Labor Force Risk & TFP Shocks

We calibrate using the monthly labor force dataset compiled by Krusell et al. (2017). They seasonally adjusted the micro CPS data for those aged at least 16, and also applied the Abowd-Zellner and margin-adjustment corrections. The resulting dataset includes the month-to-month transition probabilities across employment, unemployment, and being out of the labor force as well as the distribution of the population across these three states. While the dataset ranges from January 1978 to September 2012, we restrict our focus to the pre-September 2007 period to exclude the Great Recession.

As the underlying data include retirees, we adjust the N transition probabilities so that we are only calibrating to working-age adults, consistent with the life cycle in our model. Census data show that retirees make up 18.7% of the adult population in the United States. This share, combined with the ergodic distribution implied by the data, allows us to back out the transition probabilities for working-age N. We estimate that the probability of a working-age N remaining out of the labor force next month is 90.2%, compared to 95.6% across all adults. This adjustment implies an average working-age labor force participation rate of 81.7%, and an average employment-population ratio of 76.6%.

We assume that aggregate TFP shocks follow a three-point process. We set labor force transition probabilities in the mean TFP state to the long-run averages of the data. We calibrate transition probabilities under the low and high states by matching moments related to the employment-population ratio, unemployment rate, labor force participation rate, and off-diagonal transition rates (e.g., the probability of moving from employed to unemployed). In particular, we look at the standard deviations and autocorrelations of these variables, and the correlation with the unemployment rate of the employment-population ratio and labor force participation rate. The calibration process is over-identified, with 20 targets for 12 parameters. We employ simulated method of moments, relying on simulations of 30,000 months. As in Krusell et al. (2017), we aggregate the data to quarterly frequency, and apply the HP filter with $\lambda = 1600$. The results below are the means of 100 10,000-quarter simulations.

As Tables 3 and 4 show, across headline labor force measures, the model generally does a reasonable job fitting standard deviations, autocorrelations, and cross-correlations with the unemployment rate. However, the model struggles to capture the relatively low autocorrelation of the labor force participation rate. It also overstates the autocorrelations of transition probabilities, though it fits their standard deviations fairly well. In the Appendix, we discuss the model's fit to untargeted labor-force moments.

Table 3: Model Fit on Labor Force Metrics

	Standard Deviation		Autocorrelation		Correlation with Unemployment Rate	
	Data	Model	Data	Model	Data	Model
Employment-Population Ratio	0.0082	0.0074	0.9025	0.9127	-0.9570	-0.9586
Unemployment Rate	0.0934	0.0822	0.9111	0.8856	–	–
Labor Force Participation	0.0024	0.0033	0.6206	0.9197	-0.3002	-0.3724

Table 4: Moments on Labor-Force Transition Probabilities

	Autocorrelation		Standard Deviation	
	Data	Model	Data	Model
EU	0.5102	0.7190	0.0731	0.0761
EN	0.2952	0.7172	0.0756	0.0862
UE	0.7201	0.7186	0.0705	0.0767
UN	0.6550	0.7149	0.0957	0.0941
NE	0.3678	0.6879	0.0907	0.0920
NU	0.3063	0.5508	0.0591	0.0131

E: Employed, U: Unemployed, N: Not in labor force

3.2.2 Labor Force Risk & the Great Recession

The calibration exercise for how labor force transitions vary with TFP shocks focuses on the pre-Great Recession period. We are operating under the assumption that labor force dynamics were unusual during the Great Recession, and hence we need a different set of transition probabilities to account for this. We use our low-growth transition matrix as a starting point, and adjust this in accordance with data.

As our model is at a quarterly frequency, we rely on data from the Survey of Income and Program Participation (SIPP) to examine labor-force transition probabilities across longer time horizons. SIPP is a nationally representative survey, in which individuals are questioned once every quarter for the duration of the panel. Respondents provide information on their labor force status at a weekly frequency. We use SIPP data instead of CPS because the latter provides no more than four consecutive months of observations on any individual. Previous work has noted inconsistencies between SIPP and CPS aggregates (e.g., Mazumder, 2007; Krueger, Cramer, & Cho, 2014). In particular, SIPP is prone to "seam bias", in which labor force movements are more likely to be reported during the months in which a respondent is interviewed relative to previous months for which the respondent is reporting. We mitigate the implications of such bias for our work in two ways. First, we aggregate observations to three-month periods. As such, if a transition truly occurred in month 2, but seam bias led to it only being recorded in month 3, then our analysis would be unaffected by the error. Additionally, we compare transition probabilities across two survey cohorts. If the bias is time-invariant (e.g., due to the inherent structure of the survey), the differences between two cohorts won't be affected.

To keep our adjustments parsimonious, we targeted a single moment: the probability that someone out of the labor force remains out in the next quarter. The probabilities of such an individual becoming employed or unemployed adjust proportionately, while the transitions for those currently employed or unemployed remain unchanged. We relied on two panels of SIPP data: 2001 and 2008. These capture the Great Recession, and the most recent prior recession in the US. The sample was restricted to 25-64 year-olds, and an individual's status was recorded at the start of a quarter. We defined someone as being out of the labor force in a given month if they did not have a job at the end of the month and had not been actively looking for work during the month. This allows us to be consistent with the CPS definition of the unemployed, who must have actively sought employment recently. Over the 2001 recession, 2.32% of those out of the labor force were in the labor force a quarter later. Over the Great Recession, this entry rate was 2.16%. This is a decline in entry probability of 6.9%. We adjust the probability of remaining out of the labor force in order to match this decline.

Figure 1 illustrates the results. We examine a 60-month simulation, initialized with September 2007 data and running through September 2012. The black line shows the data, as adjusted by Krusell et al. (2017). The red line shows the results implied by the high-uncertainty transition matrix. For comparison purposes, the blue line shows the trajectories over a period of low TFP growth. Recall that the only differences between the high-uncertainty and low-TFP matrices are from the adjustment of the persistence of being out of the labor force.

Figure 1: Labor Force Trends over Great Recession, September 2007-September 2012

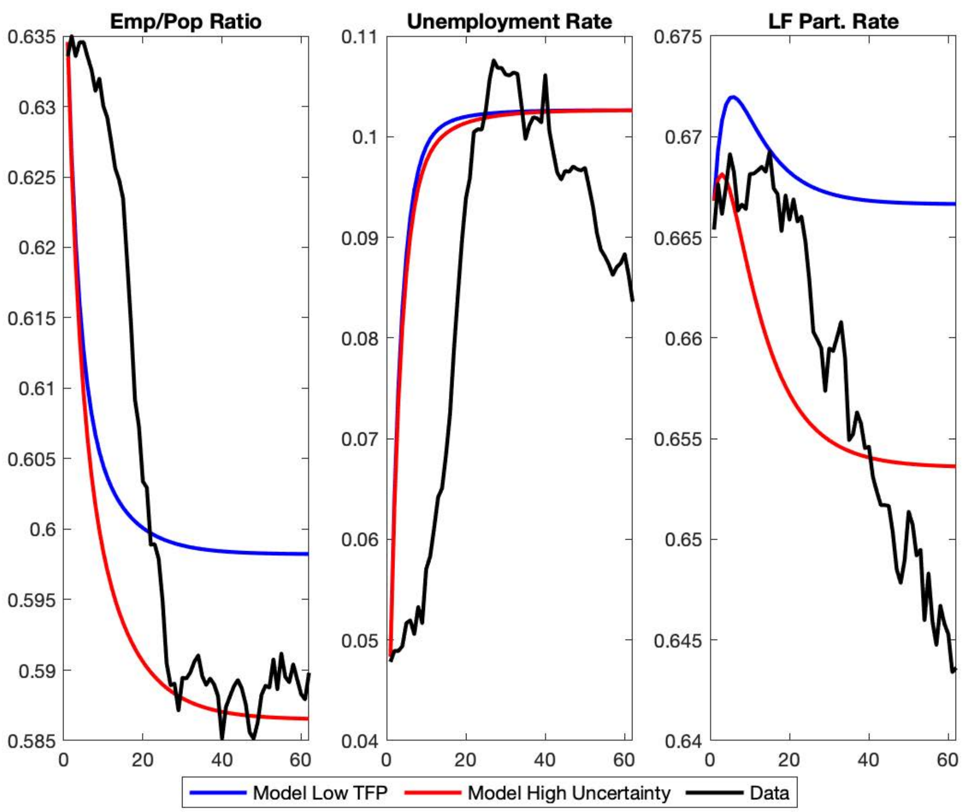

The high-uncertainty response captures the trough level of the employment-population ratio well. Its performance across the other two metrics is more mixed. While it comes close to capturing the elevated unemployment rate, the model unemployment rate remains elevated around 10%, while in the data it fell during the 2010s. These people largely exited the labor force, pushing the data measure about a percentage point lower than what the model captures. Hence, the model's implications for high-uncertainty states are relatively conservative–it understates the share of people out of the labor force and overstates the share unemployed. As the unemployed are more likely to become employed than those out of the labor force, the model will imply a swifter labor-force recovery than what we observed over the Great Recession. In that sense, we view our model as offering a conservative estimate of the severity of such a recession.

The low-TFP line in Figure 1 illustrates how different the Great Recession was from a "normal" recession. Recall that the low TFP transitions were calibrated off data prior to the Great Recession. Hence, they capture what we would expect from a recession similar to those from 1978-2007. We can see that the Great Recession was unusual not only for its duration, but for its labor force dynamics as well; a "normal" recession as long as the Great Recession would be far milder. Employment as a share of the population would be 1.5 percentage points higher, and there would be a negligible change in the labor force participation rate. The relative success of the high-uncertainty state, in terms of matching the data, indicates that an increase in the persistence of being out of the labor force can account for a lot of what made labor force dynamics during the Great Recession so unusual.

3.3 Solution

As is reflected in 1, we apply a two-step process to determine our earnings shocks. First, we the leptokurtic income shocks process is estimated to match the moments from the SSA data from Guvenen et al. (2021). Next, we discretize this process using a 25 million period simulation. We choose 15 points in its support and simulate 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 generalisation of the Backwards Induction algorithm of Reiter (2010) described by Kim (forthcoming) 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 $\mu(i,j,S,e,\zeta;z,K)$, we conjecture a value for $\widehat{K}'$, the future approximate aggregate state. Next, we determine the actual end of period aggregate capital stock by solving household value functions across the idiosyncratic state space. 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 $v^i_{n+1}(J, S, e, \zeta, z, K)$ using (3) where n indexes the iteration of value functions, across ages, with respect to the approximate aggregate state (z, K). Using $v^i_n(j+1, S', e', \zeta', z', K')$, $j = 1,..., J-1$, we solve decision rules using the endogenous grid method of Carroll (2005). Alongside $v^i_{n+1}(j, S, e, \zeta, z', K')$ this provides decision rules $S' = g^i_{n+1}(j, S, e, \zeta, z, K)$. Thus, using the proxy distribution, we can compute the 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 (1999), 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 only 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 (forthcoming) provides a feasible reduction, aggregating over (i, j, e, ζ) using weights $\omega_{i,j,e,\zeta}(S)$, 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 labor 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 evolves. 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 and employment shocks.

Table 5: Percentiles of the wealth distribution, Gini coefficient and negative assets

	1	5	10	50	90	Gini	$\leq$0
SCF	0.29	0.51	0.64	0.97	1.0	0.77	0.09
Model	0.17	0.50	0.73	0.97	1.0	0.79	0.17

Table 5 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 two 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 17 percent of the total wealth of the economy, while, in the data, the corresponding number is 29 percent. However, this overlapping generations economy, with an estimated income shock process, is able to generate an unusually large level of inequality relative to others explored in the literature.

Life-cycle wealth accumulation has an important role in generating inequality. Its effect is partly offset by precautionary savings by lower-income households. When we consider an alternative economy with no earnings risk, while retaining employment risk, the Gini coefficient rises to 0.89 as the wealth share of the poorest 90% of households falls from 27% to just 8%. If we instead consider an economy that is otherwise similar, but with expected utility with a constant relative risk aversion with $\sigma = 2$, there is less inequality. Lower levels of risk aversion reduce savings by higher-income households, leading to a fall in the Gini coefficient and a smaller share of indebted households.

Returning to the benchmark economy, with Epstein-Zin preferences and high earnings risk, Figure 2 shows the distribution of wealth for households across all ages. The left horizontal axis is the logarithm of productivity, while the right axis shows log wealth.[4] (Since some households are in debt, all wealth values are scaled up by 1 to generate the positive numbers needed to use the log scale.) Most households have relatively little wealth and the distribution of wealth across households is highly skewed. There is considerable dispersion and higher-income households tend to be wealthier given the persistence in the estimated earnings shock process. Accordingly, the dispersion is far more substantive for households with above-average productivity, whereas households with below-average productivity are disproportionately concentrated around the borrowing limit.

Figure 2: The Steady State Distribution of Wealth

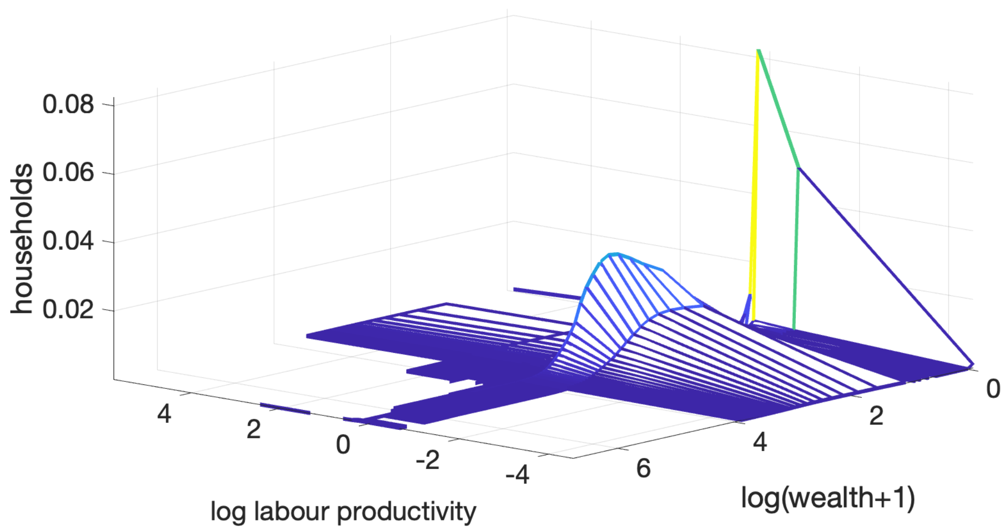

Figure 3 shows the average levels of post-tax earnings (including unemployment and retirement benefits), consumption and wealth across a generation over time. Its members retire after 45 years of work. As mentioned in the beginning of section 3, the social security replacement rate used in the model is 0.40 of average earnings, which is high. As this is independent of individual characteristics, differences in individual earnings disappear. This income, common to all retired households, is shown in the top panel starting with age 70. Relatedly, Figure 4 illustrates that variation across earnings falls to zero upon retirement.

At annual frequencies, the equilibrium real interest rate in the steady state is 5.6 percent, while households' subjective discount factors vary between 0.91 and 0.999999. Earnings risk and life-cycle savings drive the accumulation of wealth seen in the bottom panel of

[4]We continue to track the productivity of retired households, though they no longer engage in the labor market. Hence, there is a non-degenerate distribution of retirees across both dimensions in Figure 2.

Figure 3: Averages for a Cohort of Households by Age

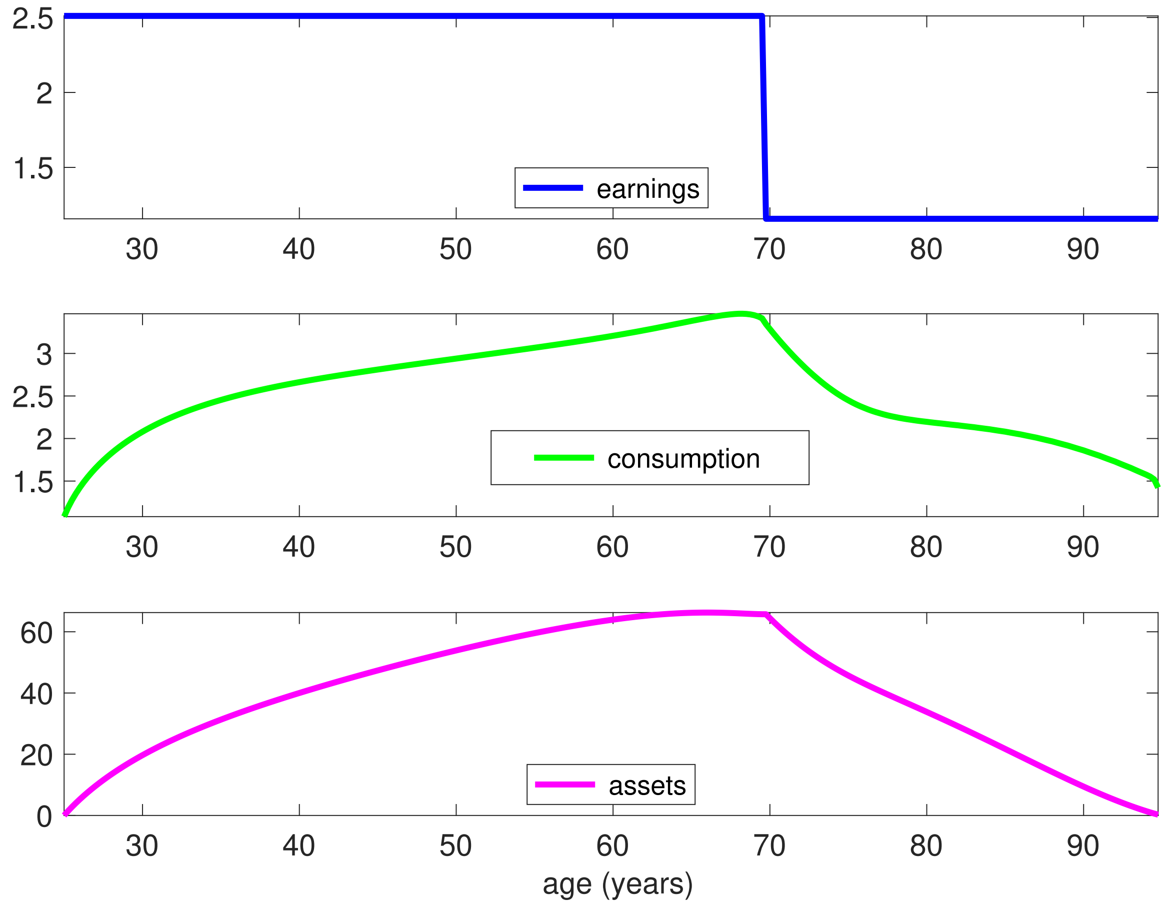

Figure 3. On average, households save until retirement. Thereafter, average wealth in the cohort falls. Survival risk, after retirement, is insufficient to fully offset the incentive to dissave given finite lifetimes. Nonetheless, most households leave accidental bequests. While working, the average rise in wealth and total income implies a gradual increase in consumption. Consumption falls gradually after retirement, as households increase their savings rate 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. Accordingly, as illustrated in Figure 4, the dispersion of assets over age follows a similar trajectory as average assets–wealth inequality is highest just prior to retirement, but subsequently falls as uncertainty over lifespan dissipates.

The inequality generated by the model, solely from earnings and employment risk propagated by life-cycle savings, suggests that it is a useful starting point to evaluate the roles

Figure 4: Standard Deviations for a Cohort of Households by Age

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

4.2 Business Cycles

We study business cycles driven by shocks to total factor productivity growth. As explained above, these growth shocks alter households' labor market risk. Hence we follow an approach in İmrohoroğlu (1989) and Krusell and Smith (1998) of having exogenous changes in employment over the business cycle. As in Krusell et al. (2017), we assume three labor market states with individuals either employed, unemployed or out of the labor force. Since the probabilities of moving from one state to another vary with TFP growth, there are aggregate changes in the distribution of uninsurable risk across individuals. In recessions, the chance of being out of the labor force rises. This, alongside cyclical variation in the probability of unemployment, drives fluctuations in the number of employed workers over the business cycle.

Table 6: Business Cycle Moments

	Z	Y	C	I	N	K	w
std	2.29	1.84	1.25	4.40	0.97	0.94	1.45
correlation	0.88	1.00	0.81	0.90	0.61	0.14	0.85

Note: All series are HP-filtered with a weight of 1600. The series are TFP (Z), GDP (Y), aggregate consumption (C), aggregate investment (I), labor (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.

Table 6 reports business cycle moments from a 2000 period simulation of the benchmark model. This incomplete markets economy 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, as seen in figure 5, growth shocks change the mechanics of TFP. As the growth shock is persistent, and involves a period of negative growth rates, the level of TFP has a u-shaped response. This implies that the marginal product of capital is itself u-shaped, mirroring TFP. In contrast, in a model with stationary TFP, the real interest rate falls then monotonically returns to its average level, following the monotone path of TFP itself. Consequently, TFP shocks' wealth and substitution effects on households vary markedly across the two approaches.

In the typical business cycle model, the sharp, immediate fall in the real interest rate reduces the relative price of consumption goods, creating a large substitution effect at the onset of a recession. Dampening the wealth effect, this leads to a gradual, non-monotone response in aggregate consumption. In contrast, the stochastic growth and non-stationary TFP within our setting generates non-monotone paths of TFP and interest rates. A gradual decline in the marginal product of capital implies a smaller substitution effect than in a stationary TFP model. Most of the consumption response is determined by the large wealth effect resulting from permanent shocks to TFP. As a result, as in figure 6, consumption responses, unlike TFP, are monotone. This leads to the far larger variability in consumption than seen in a typical business cycle model (Table 6).

Second, employment is non-monotone and lags GDP, as it does in the data. In Table 6, this is seen in a lower contemporaneous correlation between Y and N. Aggregate employ-

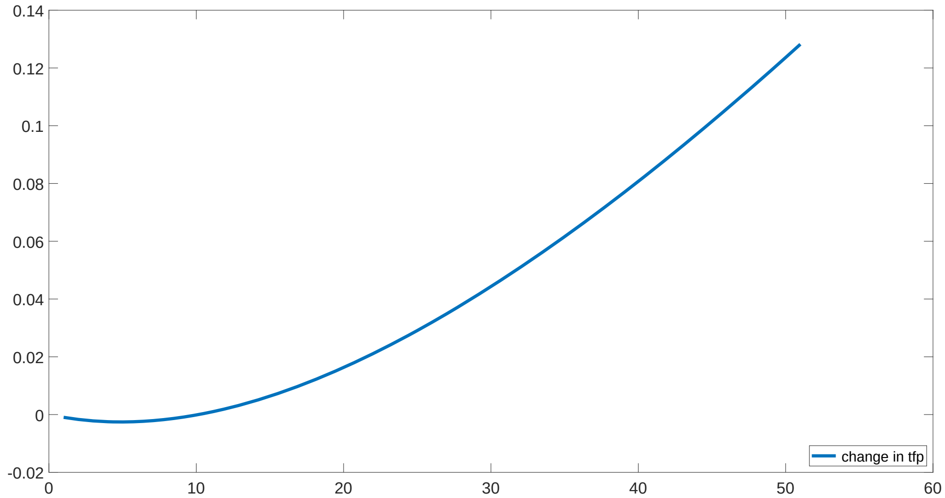

Figure 5: Total Factor Productivity following a persistent growth shock

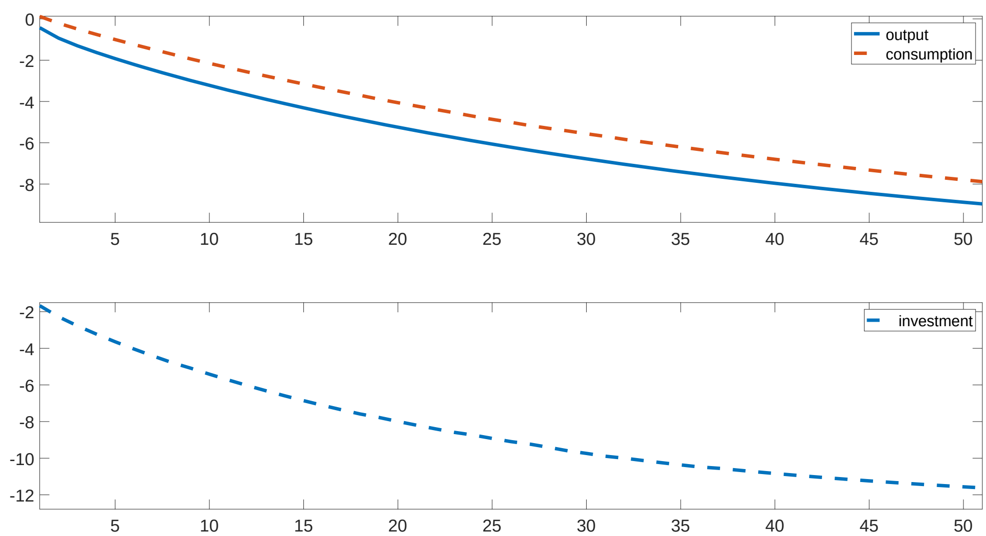

Figure 6: GDP, Consumption and Investment

ment is determined by the distribution of employment across households. When transition probabilities over employment states change with TFP growth, the distribution of households responds gradually. At the level of the individual household, the gradual response is beneficial at the onset of a recession; even in recessions, employment is a highly persistent state. However, this gradualism also means that the recovery of the labor force lags the recovery in the production sector; most non-employed households will experience multiple additional periods of joblessness before becoming employed. Consequently, recessions and jobless events entail longer-lived risks to individual workers than the economy as a whole, which further motivates precautionary savings and accounts for the monotone responses in aggregate consumption. Both the lower correlation between consumption and GDP, and the lag in employment, move the incomplete markets model business cycle closer to the data, compared to the standard expected utility, representative agent model. However the gradual response of employment to a TFP shock implies lower variability relative to GDP. The slower response in employment reduces the correlation between the real wage and GDP which, at 0.85, is smaller than that usually seen in a model with variable labor supply.

4.3 The Great Recession and Lost Recovery

We now turn to the Great Recession and the decade that followed. We argue that changes in TFP, by themselves, cannot explain the aggregate dynamics observed over this period. GDP, consumption and investment remained depressed through 2018 (see figure 7), while TFP quickly returned to its trend path. Given the persistent fall in labor force participation discussed above, we turn to changes in labor market risk and explore the extent to which they help explain the ten years after the Great Recession.

Starting with the Great Recession itself, Table 7 reports the percent changes in all series, except TFP, between the 2007Q4, the start of the recession, and 2009Q2, its trough. The change in TFP is reported between 2007Q4 and 2009Q1, when it reached its lowest level. The recession saw GDP fall by 9 percent, relative to its long-run trend. Consumption, similarly detrended, fell 6.6 percent over the same period while private investment dropped by 29 percent and total hours worked fell 9.5 percent. Surprisingly, measured TFP fell by only 3.3 percent.

Table 7: The Great Recession: Data and Model

	TFP	GDP	C	I	N
07Q4 to 09Q2	-3.26	-8.96	-6.55	-29.02	-9.53
model	-3.26	-4.24	-2.46	-8.04	-4.64

The data series are linearly detrended using their long-run growth rates averaged over 1954Q1 and 2007Q4. We report each series as a percent deviation from its 2007Q4 value.

We explore the extent to which our model economy, with empirically consistent shocks to TFP growth is able to reproduce the large fall in GDP and consumption observed over the recession. We choose shocks to the growth rate of TFP to imply a 3.26 percent fall over five quarters, then a return to trend over the following four quarters. This matches the path of TFP over the recession, seen in Figure 7. We begin the recession with a growth rate shock of -0.88 percent. There is a second positive shock of the same magnitude 5 quarters later. We assume that the persistence of the realized shocks is 0.9605. As the initial growth shock is persistent, the level of TFP continues to fall until the second, positive shock initiates the recovery. Lastly, 8 quarters after the beginning of the recession we return quarterly TFP growth to its mean value of 0.4 percent.

There is a second force driving the model recession, a rise in labor market risk. This not only increases uninsurable income risk for households; it also gradually reduces aggregate employment. Importantly, households' employment risk follows the transition probabilities we estimated using data specific to the recession period. Specifically, we find that the persistence of remaining outside the labor force was unusually high over the Great Recession. Accounting for this leads to longer durations of non-employment at the household level. At the level of the aggregate economy, it also implies a persistent fall in total employment, leading to a larger decline in total hours than would be seen in an ordinary recession.

The second row of Table 7 reports the peak to trough changes in our model's aggregate series. All variables are reported in terms of percent changes over 6 quarters, except for TFP, which matches the 5-quarter change used for the data. Our incomplete markets business cycle model is able to explain almost half, 47 percent, of the peak-to-trough decline in output. Further, it reproduces 38 percent of the fall in consumption, and 49 percent of the reduction in total hours, seen in the data. However, in the absence of financial shocks, the fall is investment is less pronounced, at 28 percent of the data. The drop in investment is dampened by the strong precautionary savings motive, which contributes to a fall in consumption and a rise in savings as households respond to the rise in their earnings risk.

As seen in Figure 7, the Great Recession saw a relatively brief fall in TFP growth alongside a very protracted decline in the real economy, relative to long-run trends. We now explore how much of the trend in hours worked over the decade after the Great Recession can be captured by a persistent rise in labor market risk when TFP growth returns to trend after 9 quarters.[5] The path of TFP growth, which reproduces the peak to trough fall and recovery seen in Figure 7 is shown in Figure 8 in blue. The red series shows the fall in aggregate hours worked as a result of the rise in labor market risk that we

[5]Recall that our increased labor market risk solely stems from an observed change in the persistent of being out of the labor force during the Great Recession. Hence, it only indirectly affects total employment and hours worked.

Figure 7: Ten Years Gone: After the Great Recession

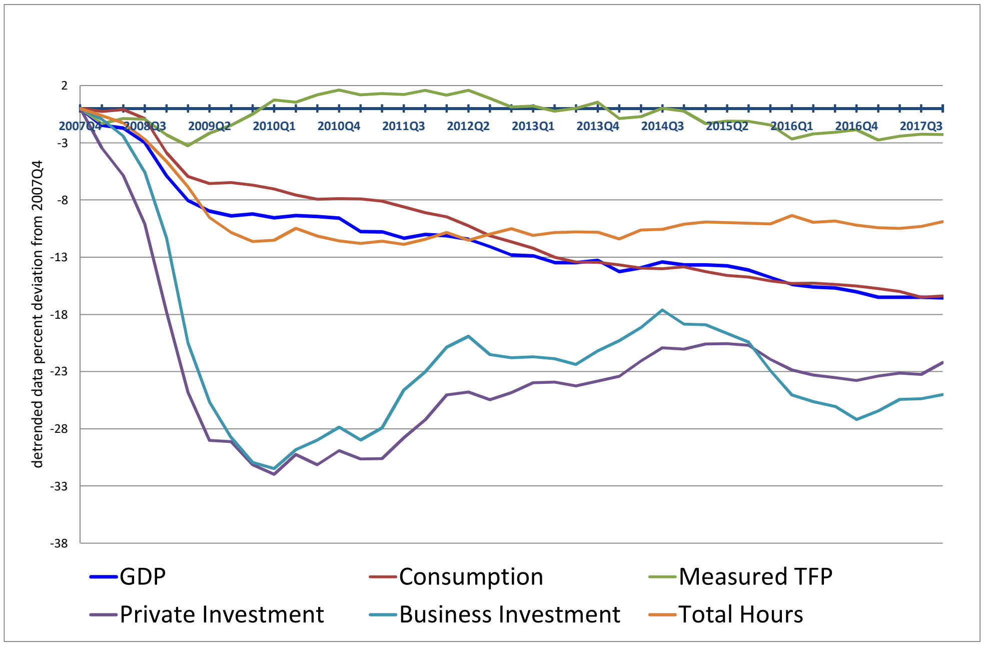

Data sources: total hours [Cociuba et al 2018]; other data: BEA May 2018

have estimated over the Great Recession. Note the gradual decay in aggregate labor as the distribution of households responds over time to the change in uninsurable risk. However, our estimated rise in risk does not fully explain the fall in total hours seen in Figure 7. By the third quarter of 2011, total hours worked had fallen 11.9 percent and, as late as the end of 2017, it remained 9.9 percent below its trend. Hence our estimated labor market transition process explains 52.4 percent of the observed decline in the series.

The decline in hours worked in the model recession is dampened by the relatively mild fall in aggregate investment, which in turn implies a smaller decline in the marginal product of labor. This is reflected in the trends shown in Table 7; the model captures a larger share of the fall of consumption than of investment seen in the data. The rise in uninsurable income risk over the recession results in a sharp increase in precautionary savings by households. In our closed economy, stronger precautionary savings imply a dampened fall in aggregate investment. In the absence of an investment-specific shock reducing the productivity of investment spending, there is only a mild fall in aggregate capital, despite the 5.2 percent fall in aggregate labor. In the model, capital was 44.25 before the recession and falls to 41.33

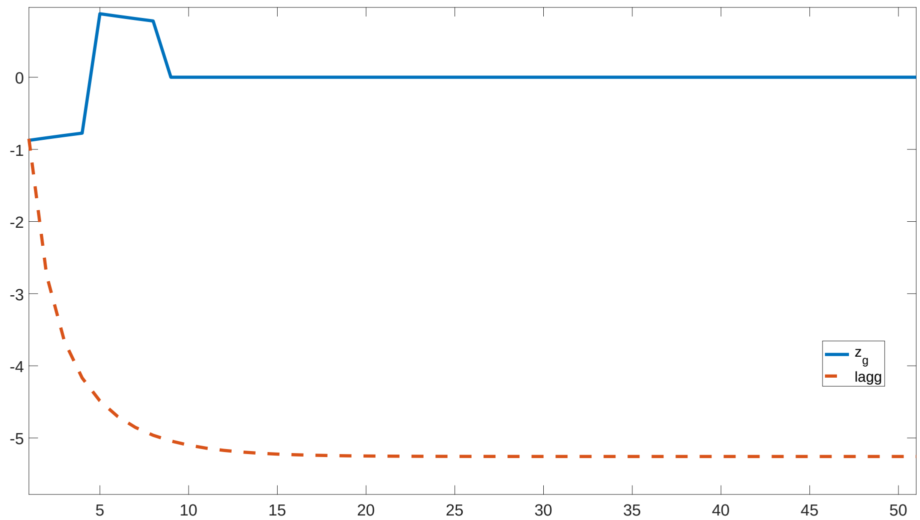

Figure 8: TFP growth and Employment

ten years later. Thus precautionary savings raises the capital to output ratio, dampening the effect of the decrease in labor. Overall investment falls only 4 percent over the decade, and output 4.5 percent. As a result, consumption declines by 2 percent over that span.

In the data, we see that aggregate investment fell 22.2 percent through 2017Q4. To avoid a counterfactually small decline in investment, we now assume an exogenous shock to the productivity of investment spending.[6] This reduction in investment productivity, which remains in effect throughout the period examined, helps move the fall in investment in the model closer to the data. It also partly offsets the strong precautionary savings motive.

In Figure 10 we now see a large effect of the rise in labor market risk when coupled with a 20 percent decrease in the productivity of investment spending. Output and consumption eventually fall 10.7 and 9.7 percent, respectively. There is a corresponding 7.3 percent fall in investment through the decade after the recession's start. Interestingly, the initial decline in investment is similar to that seen in the data, but in later periods households' attempt to self-insure against the rise in income risk leads to a recovery in investment, even as consumption and output continue to fall. This echoes the trends observed in the data (Figure 7).

[6]The financial crisis that prompted the Great Recession disrupted intermediation. An exogenous shock to investment productivity is a reduced-form way of capturing this.

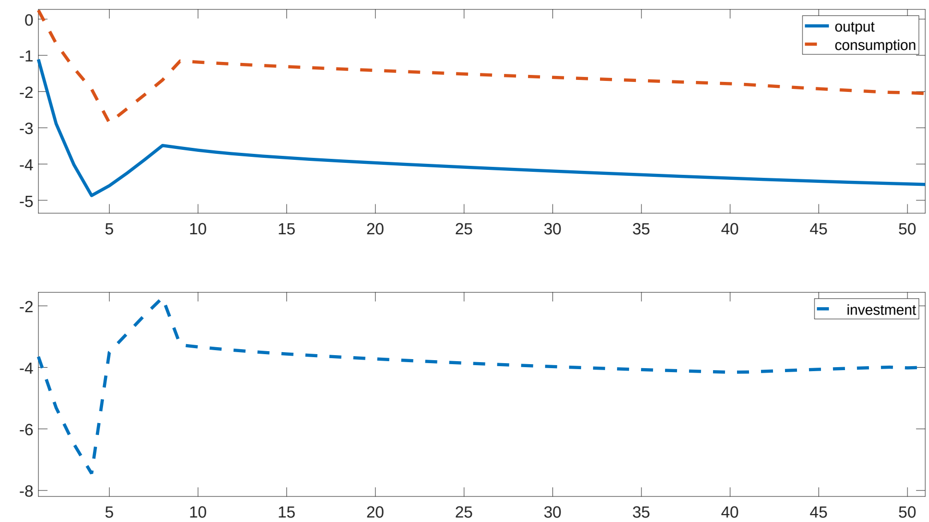

Figure 9: GDP, Consumption and Investment

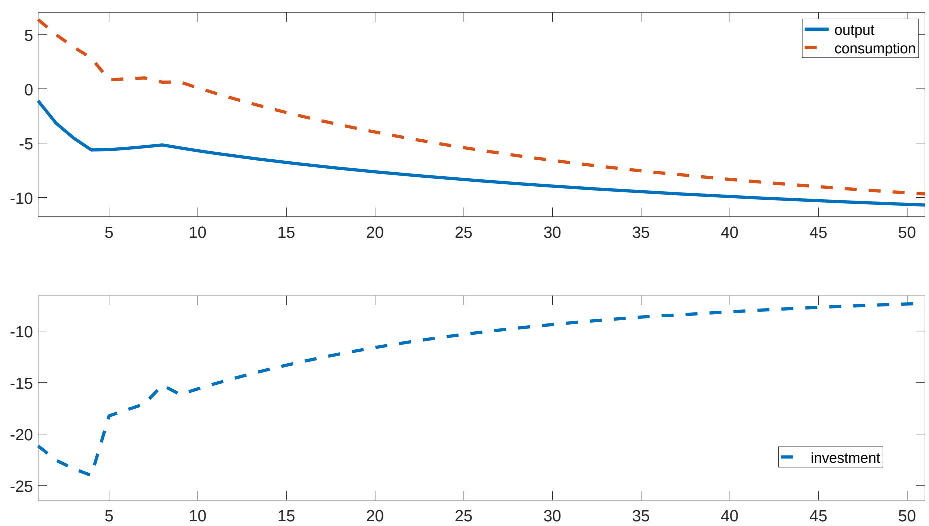

Figure 10: GDP, Consumption and Investment with Investment Shock

Thus our incomplete markets economy can produce a persistent recession in aggregate quantities that is consistent with the absence of a persistent reduction in TFP. The lack of economic recovery is driven by an estimated rise in labor market risk. However, this rise in risk strengthens households' precautionary savings motives, which provide upward support to aggregate capital. This makes the large, persistent fall in aggregate investment over the Great Recession all the more remarkable. We show that a reduction in the efficiency of investment spending can help close the gap between model and data. Overall the model reproduces 60 percent of the persistent decline in output and consumption, and one-third of the fall in private investment. It generates 52.4 percent of the decline in total hours worked, hence the relatively large decreases in GDP and consumption are notable. A larger shock to investment spending would result in larger falls in these series.

5 Concluding Remarks

We have developed a quantitative overlapping generations model where households face both earnings and employment risk. These two large sources of income risk, calibrated to match the data, 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 need to save for retirement.

Finite lifetimes and life-cycle savings motives imply large individual consumption responses over the business cycle. Aggregate consumption is more volatile than in a model with infinitely-lived households. This volatility in aggregate consumption is reinforced by the non-monotone response in aggregate total factor productivity that follows from a persistent shock to growth rates. The resulting non-monotinicity of the real interest rate implies a large monotone response in aggregate consumption. Moreover, a rise in labor market risk, consistent with the fall in labor force participation following the Great Recession, increases precautionary savings and contributes to a persistent fall in consumption.

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,' *Economics Letters*, vol. 91, no. 3, pages 312-320.

[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 Jos-Vctor Ros-Rull (2020) 'Intergenerational Redistribution in the Great Recession,' *Journal of Political Economy*, vol. 128, no. 10, pages 3730-78.

[6] Guerrieri, Veronica and Guido Lorenzoni (2017) 'Credit Crises, Precautionary Savings, and the Liquidity Trap,' *Quarterly Journal of Economics*, vol. 132, no. 3, pages 1427-67.

[7] Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, Jae Song (2019) 'What do Data on Millions of US Workers Reveal about Life-Cycle Earnings Dynamics?' *Econometrica*, vol. 89, no. 5, pages 2303-39.

[8] Hansen, Gary (1985) 'Indivisible Labor and the Business Cycle,' *Journal of Monetary Economics*, vol. 16, pages 309-27.

[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] İmrohoroğlu, Ayşe (1989) 'Cost of Business Cycles with Indivisibilities and Liquidity Constraints.' *Journal of Political Economy* vol. 97, no. 6, pages 1364-83

[12] Kaplan, Gregory, Benjamin Moll and Giovanni Violante (2018) 'Monetary Policy According to HANK' *American Economic Review* vol. 108, no. 3, pages 697-743.

[13] Kim, Heejeong (forthcoming) 'Inequality, Portfolio Choice and the Business Cycle' *Review of Economic Dynamics*

[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, Alan, Judd Cramer, and David Cho (2014) 'Are the Long-Term Unemployed on the Margins of the Labor Market?' *Brookings Papers on Economic Activity* Spring 2014: pages 229 – 299.

[16] 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.

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

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

[19] Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin (2017) 'Gross Worker Flows over the Business Cycle' *American Economic Review*, 107(11), pages 3447 – 3476

[20] Mazumder, Bhashkar. (2007) 'New Evidence on Labor Market Dynamics over the Business Cycle' *Federal Reserve Bank of Chicago Economic Perspectives*, 1Q/2007, pages 36 – 46

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

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

Appendix

A1: Additional Moments on the Labor Force Transition Matrices

To illustrate labor force dynamics under our calibrated transition matrices, we consider counterfactual scenarios in which an economy is permanently subject to a given set of transition probabilities. Then the stationary distributions of working-age individuals would be

Table A1: Distribution of Working-Age Population under Permanent Processes

	E	U	N
Low-Growth	0.7368	0.0843	0.1789
Medium-Growth	0.7671	0.0513	0.1815
High-Growth	0.8128	0.0326	0.1546

E: employed, U: unemployed, N: not in labor force

These stationary distributions imply the following metrics across the full population.

Table A2: Aggregate Labor Force Metrics under Permanent Processes

	Emp/Pop	Unemployment	Labor Force	Unemployment Duration	
	Ratio	Rate	Part. Rate	Mean	Median
Low-Growth	59.9	10.3	66.8	10.89	2.17
Medium-Growth	62.4	6.3	66.5	2.65	1.52
High-Growth	66.1	3.9	68.7	2.08	1.08

Emp/Pop: Employment/Population

Part.: Participation

Duration measured in months

These labor force statistics give some intuition on what an especially prolonged expansion and contraction would look like in the model. Against this benchmark, both the employment-population ratio and the unemployment rate look qualitatively reasonable relative to the sample period (1978-2007). However, there is very little movement in the labor force participation rate. Moreover, the calibration implies that the labor force participation rate is lowest in the medium state, rising in both downturns and upturns. The average duration of unemployment during low-growth states (10.9 months) is also rather high, though the much milder rise in the median duration indicates that for most workers the realized increase in risk is not nearly as large.

Untargeted Moments on Labor Force

Table A3 examines the model's labor force metrics across moments not targeted in the calibration process. Specifically, it focuses on the first and third quartiles of the data, as well as the series' skewness. In general, the model is able to capture these quartiles fairly well, particularly for the employment-population ratio. The largest discrepancy is in the 25th percentile of the labor force participation rate, with the model unable to generate as large a decline as in the data. The employment-population ratio accordingly offsets this, and is shifted somewhat higher in the model than in the data. More broadly, the model is less effective in capturing the skewness in the data, with each measure being close to 0. This contrasts with the strong negative skewness of the employment-population ratio in the data, and the positive skewness of the unemployment rate.

Skewness is imprecisely identified: all measures are within 1.5 standard deviations of 0.

Table A3: Comparison of Labor Force Metrics on Untargeted Moments

	25th Percentile		75th Percentile		Skewness	
	Data	Model	Data	Model	Data	Model
Employment/Population	0.607	0.614	0.636	0.645	-0.675	0.110
Unemployment Rate	0.055	0.040	0.075	0.068	0.382	-0.008
Labor Force Participation Rate	0.634	0.665	0.674	0.679	-0.155	0.057

Data: Quarterly averages from 1978Q1 through 2007Q4, using Krusell et a. (2017) dataset.

Model: Moments taken across 100 10,000-quarter simulations.