Countercyclical Earning Risk in Large Recessions: Consumption Dynamics and Welfare Losses*

Bernardo Candia† Francisco Díaz-Valdés‡

UC Berkeley Universidad de Chile

September 19, 2022

Preliminary

Abstract

Labor earnings exhibit a pro-cyclical skewness. We investigate how idiosyncratic cyclical risks account for the Great Recession's consumption dynamics and assess how these risks exacerbate welfare losses. Building a model of incomplete markets with idiosyncratic unemployment and efficiency risk correlated with aggregate shocks, we find three major differences when comparing it to a model with only cyclical unemployment. First, consumption's initial decline is 0.5 percentage points larger, and its recovery is slower. Second, aggregate welfare losses are bigger in terms of lifetime consumption: 4.1% versus 3.1%. Third, the cross-sectional distribution of household welfare losses has a thicker and longer right tail, meaning a non-negligible fraction of households suffer major losses.

Keywords: Heterogeneous Agents, Idiosyncratic Income Risk, Cyclical Income Risk, Large Recessions, Welfare.

JEL: D31, E21, E32, J31.

*We thank Álvaro Aguirre, Valerie Boctor, Tomas Breach, Eduardo Engel, Álvaro García, Nick Gebbia, Yuriy Gorodnichenko, and Oliver Kim for helpful comments and suggestions.

†bernardo_candia@berkeley.edu

‡fdiazvalde@fen.uchile.cl

1 Introduction

Recent literature on labor earning dynamics describes recessions as times when employed households face higher downside risks to their current and future labor income. For instance, Guvenen, Ozkan, and Song (2014) and Busch, Domeij, Guvenen, and Madera (2022) find that the distribution of earnings growth displays substantial pro-cyclical skewness. During recessions, large upward earnings movements become less likely, whereas large drops in earnings become more likely.[1] Furthermore, the literature has documented that these labor earnings risks seem to be highly persistent.

Motivated by this new empirical evidence, we build on Krueger, Mitman, and Perri (2016a) (henceforth KMP) and McKay (2017) to develop a real business cycle model that features heterogeneous households, incomplete markets, and idiosyncratic earnings risk that is correlated with aggregate shocks. Using our model, we tackle three main questions. (i) How much do household wealth inequality and idiosyncratic cyclical labor earnings risk account for the initial response of consumption to large real aggregate shocks? (ii) How do these cyclical risks shape the recovery of consumption after the initial drop? (iii) How much do these cyclical risks exacerbate aggregate welfare losses from severe recessions, and how do they shape the cross-sectional distribution of household welfare losses?

In our model, the cyclicality of labor income risks comes from two sources. First, in the spirit of Krusell and Smith (1998), unemployment is exogenously determined, and its duration and persistence increase during recessions. Second, conditional on employment, the distribution of labor earning income risk displays pro-cyclical skewness. Because the literature still does not provide a widely-accepted theory on why the distribution of employed households' earnings risk exhibits cyclical skewness, we follow McKay (2017) and McKay and Reis (2021) taking as given these cyclical changes. We then proceed to analyze its consequences for consumption dynamics and welfare.

[1] See the Global Repository of Income Dynamics website for an extensive list of current research on income dynamics for various countries.

To illustrate the effects of these time-varying earnings risks, we report our findings relative to a model where only unemployment risk varies over the cycle.

We start our investigation by exploring if the inclusion of countercyclical earnings risk can affect the ability of our model to reproduce the main empirical facts characterizing the US wealth distribution. Evaluating the effects of the addition of cyclical income risk is crucial because KMP has shown that incomplete markets must feature a realistic wealth distribution to produce large consumption drops in response to real aggregate shocks. We find that incorporating the cyclical nature of earnings risk does not significantly alter the ability of our model to replicate the US wealth distribution, to the contrary, our model does a great job matching the observed US wealth distribution. For instance, the mean square distance of the share of net worth held by quintiles between data and the model is 1.8, while the mean square distance between data and the KMP model is 6.9. Moreover, the model generates wealth-poor households that represent a larger share of aggregate consumption than in KMP, which is vital to greater consumption responses.

Then, we perform two experiments to evaluate the impulse responses of macroeconomic aggregates to negative technology shocks, focusing mainly on the response of aggregate consumption. The two experiments are (i) a one-time negative technology shock and (ii) a Great Recession type shock that lasts, on average, 22 quarters. In the first experiment, when the recession hits the economy, the initial decline in consumption is 0.5 percentage points larger than the economy with just cyclical unemployment risk. The larger consumption drop is because persistent earnings decline becomes more likely, increasing the expected duration of the earnings losses during economic meltdowns. In the second experiment, we find the initial drop is more persistent, languishing the consumption recovery compared to the economy with just cyclical unemployment. The reason is that the worst economic outlook leads households to increase their precautionary savings, generating a substantial reduction in consumption and a slow recovery. In evaluating the consumption response to severe recessions, we differ from McKay (2017) in an important aspect as we consider the role of the negative productivity shock in reducing factor prices and decreasing households' disposable income even further in severe recessions.[2]

Finally, following Krueger, Mitman, and Perri (2016b), we measure the welfare losses from experiencing severe economic meltdowns such as the Great Recession and study how they are distributed across the population. In the model with cyclical labor income risks, aggregate welfare losses are around 4.1% of lifetime consumption, representing an increase of one percentage point in welfare losses compared to the model with just unemployment risk. Moreover, the model with both types of risks displays a welfare loss distribution with a thicker and longer right tail, meaning that a non-negligible fraction of households suffers major losses. For instance, around 23% of households experience losses bigger than 5% of lifetime consumption, while it is about 11% in the model with just cyclical unemployment.

Related literature. Since Krusell and Smith (1998) influential paper, understanding the role of incomplete markets and household heterogeneity in the business cycle has become an active area of research. This paper adds to the growing literature on the relationship between wealth inequality and real macroeconomic shocks. Krueger, Mitman, and Perri (2016a), which is our most related paper, studies an incomplete markets model with idiosyncratic income risk and preference heterogeneity to quantify how household heterogeneity, particularly wealth inequality, amplifies and propagates negative aggregate shocks. Their key finding is that net worth inequality significantly deepens the aggregate consumption drop in response to a negative macroeconomic shock relative to the standard representative agent economy. Wealth-poor and borrowing-constrained households, which have a high marginal propensity to consume, sharply cut their consumption expenditures to increase precautionary savings as recession hits.

While the findings of KMP provide insights into how the presence of a significant fraction of households with little or no wealth exacerbates the response of consump-

[2] McKay (2017) holds total factor productivity constant in his main result. Thus, his model generates slight output declines, producing, in turn, insignificant reductions in factor prices.

tion in recessions, Amromin, De Nardi, and Schulze (2018) argues that KMP could understate the consumption drop and the subsequent weak recovery seen in the data, by abstracting from relevant changes that occurred during the Great Recession. In particular, and more important for our purposes, the KMP model assumes that unemployment is the only cyclical idiosyncratic risk while there is a vast literature documenting that in recessions, conditional on being employed, the likelihood of large and persistent earning declines increases, whereas it decreases for upward earning movements.

Guvenen, Ozkan, and Song (2014) using labor earning data from the US Social Security Administration documented that, contrary to past research supporting countercyclical variance (Storesletten, Telmer, and Yaron, 2004b), the idiosyncratic shock earning variances are not countercyclical. Rather, conditional on employment, the cyclical component comes from changes in the skewness. In recessions, the right tail of earnings shock distribution collapses while the left tail enlarges, yet the median slightly varies relative to the tails. Similarly, Busch, Domeij, Guvenen, and Madera (2022) employing administrative data from the United States, Germany, and Sweden, found that skewness is robustly pro-cyclical. Changes in hours and wages are essential to generate the pro-cyclical skewness in earnings growth. The finding of strongly pro-cyclical skewness in earnings growth conditional has also been found in the UK (Angelopoulus, Lazarakis, and Malley, 2019) and in Denmark (Harmerberg and Sievertsen, 2021). Moreover, Nakajima and Smirnyagin (2019), Busch and Ludwig (2021), Guvenen, Mckay, and Ryan (2022) have documented the same pattern using the US Panel Study of Income Dynamics (PSID).

If households expect declines in labor earnings to be long-lasting in economic downturns, they will cut consumption much more for precautionary reasons, causing aggregate consumption to fall even further. Moreover, after the onset of the recession, households would begin to form or increase their buffer stocks, weakening the consumption recovery. McKay (2017) found that a significant part of the decline of aggregate consumption during the Great Recession could be explained by the increase in the downside risks on labor earning prospects. Nevertheless, his work does not mean to be a complete depiction of the Great Recession, as he maintains constant the total factor productivity. He neither investigates the welfare losses of experiencing a severe recession nor how they are distributed across households.

Our work also relates to the literature on the welfare costs of aggregate fluctuations. In a seminal article, Lucas (1987) calculated that the gains from eliminating business cycle fluctuations are insignificant (around one-tenth of one percent of annual consumption for the US). An extensive literature has questioned the assumptions underlying his contentious result, namely: complete markets, the lack of interaction between aggregate and idiosyncratic shocks, preferences, and the use of infinitely-lived households (see Imrohoroglu (2008) for a survey). In particular, as wealth is unequally distributed, it is reasonable to presume that welfare losses will be unevenly distributed. Households at the borrowing constraint or those with little wealth cannot insure themselves from a negative income shock to the same degree as rich-wealth households.

The importance of incomplete markets and the relationship between idiosyncratic and aggregate shocks for welfare analysis was documented by Krusell, Mukoyama, Sahin, and Smith (2016). Using an incomplete markets model with stochastic discount factors and unemployment shocks, they calculated that welfare gains are around one order of magnitude larger than those computed by Lucas (1987). Likewise, and especially relevant to our work, Krueger, Mitman, and Perri (2016b) calculate the welfare losses from the Great Recession using the KMP model. They found that the welfare cost of losing one's job at the onset of the recession is 2% of lifetime consumption for the wealthiest quintile, whereas it is 5% for the poorest. The latter result is consistent with Chatterjee and Corbae (2007), as they have shown that the welfare gains from eliminating the probability of a severe deep recession, such as the Great Depression, range between 1% and 7% of lifetime consumption.

Furthermore, the nature of labor income risk is relevant to measuring the welfare costs of macroeconomic instability. Guvenen (2007, 2009), Krebs (2007), Heathcote, Storesletten, and Violante (2009), Low, Meghir, and Pistaferri (2010), and McKay and Reis (2021) all explore the welfare implications of different ways of modeling idiosyncratic income risk. We contribute to this literature by computing the welfare losses of experiencing a recession and how they are distributed across the population. We consider the relationship between the earning risk and the business cycle, and given its high persistence, it will have important welfare implications for the economy, especially for those households near or at the borrowing constraint.

This paper is organized as follows. Section 2 develops a real business cycle model with heterogeneous households, incomplete markets, and countercyclical earning risk. Section 3.2 describes the calibration. In section 4, we study to which extent the model can match the relevant features of the observed US wealth distribution. We then analyze the response of aggregate macroeconomic variables and welfare losses when the economy goes through a severe economic downturn. Section 5 concludes, and the appendix contains a detailed description of the estimation of the stochastic process for labor earnings, solution methods, complementary theory, and the computational algorithm employed.

2 Model

This section builds a dynamic general equilibrium model based on Krueger, Mitman, and Perri (2016a). The model features heterogeneous households, incomplete markets, aggregate productivity shocks, and idiosyncratic risk in the form of unemployment and labor productivity (or efficiency shocks, for the lack of a better term). The model's key feature is that idiosyncratic labor productivity shocks vary over the business cycle. As far as our knowledge is concerned, there is no well-established theoretical foundation for the cyclicality of long-term earnings changes. Therefore, we follow McKay (2017) and McKay and Reis (2021), assuming this reduced-form approach in which idiosyncratic labor efficiency varies over the business cycle, generating the procyclical skewness of labor income.

2.1 Technology

A unique final good Y is produced out of capital K and labor L by a representative firm according to a Cobb-Douglas production function:

$$Y = zf(K, L) = zK^{\alpha}L^{1-\alpha}, \quad 0 < \alpha < 1,$$

where z is an exogenous total factor productivity shock (TFP), which follows a first-order Markov chain with transition matrix $\pi(z, z')$, and the TFP shock takes values $z \in \mathcal{Z}$. We assume that $\mathcal{Z} = \{z_l, z_h\}$, where $z_l < 1 < z_h$, interpreting z_l as a severe recession and z_h as normal times. Let $\Pi(z)$ be the invariant distribution of the TFP shock. As usual, the firm maximizes profits by solving a static problem. It rents capital and labor at prices r and w, respectively, so that the following first-order conditions hold:

$$r = zf_K(K, L),$$
$$w = zf_L(K, L).$$

2.2 Households

2.2.1 Households endowments, preferences, and savings

A unit mass of households populates the economy. Households have stochastic life horizons due to a constant probability of dying in each period equal to $1 - \theta \in (0, 1)$. The fraction of deceased households is replaced by an equivalent measure of new-borns, leaving the population size unchanged.

Households derive utility from the consumption of the final good according to a CRRA utility function with relative risk aversion parameter σ. Households seek to maximize their lifetime utility given by:

$$\mathcal{W} \equiv \mathbb{E}_0 \left[\sum_{t=0}^{\infty} (\beta\theta)^t \frac{c_t^{1-\sigma}}{1-\sigma} \right],$$

where c_t is the household's consumption in period t, and β is the intertemporal discount factor, which is heterogeneous across households but fixed over time for a given household. Following Carroll, Slacalek, Tokuka, and White (2017), households draw β at the beginning of their life from a uniform distribution with support $\left[\overline{\beta} - \nu, \overline{\beta} + \nu\right]$.[3]

In each period, households have an endowment of one unit of time and a stochastic log-labor efficiency $\gamma \in \mathcal{Y}$. Households supply inelastically their unit of time with labor efficiency equal to $\exp(\gamma)$ to the labor market. Additionally, they could be either unemployed or employed. Let $\varepsilon \in \mathcal{E} = \{0, 1\}$ denote the current employment status of a household, with zero and one denoting unemployment and employment, respectively.[4] Employed households receive a pre-tax labor income equal to $w \exp(\gamma)$. In contrast, the unemployed receive an amount of b from an unemployment insurance system. The amount b is equivalent to a fraction $\rho \in (0, 1)$ of their potential labor income.[5] Following Krueger, Mitman, and Perri (2016a), we assume that taxes are levied on both labor earnings and unemployment benefits at rate $\tau(z, \rho) \in (0, 1)$, which may depend on the aggregate state of the economy.

Households can save (but not borrow) by accumulating physical capital and having access to perfect annuity markets.[6] Hence, the gross return of savings, conditional on survival, equals $(1 - \delta + r)/\theta$.[7] We denote by $a \in [0, \infty)$ the household's capital or asset holdings. In each period capital depreciates at a rate $\delta \in (0, 1)$. Since households cannot borrow, markets are incomplete. Therefore, there are no financial instruments with which households can fully insure themselves against idiosyncratic risks. Consequently, households will try to hedge by holding physical capital.

[3]With permanent discount factor heterogeneity the wealth distribution could be unbounded. However, it is not the case in the present work because of the positive probability of dying.

[4]For simplicity, we assume that the employment status is stochastic to represent, in a reduced form, the underlying frictions in the labor market.

[5]We assume that employment status and labor efficiency are public information, so only unemployed households will receive the unemployment benefits.

[6]The assumption of exogenous borrowing constraints represents the underlying frictions that households face in financial markets. While the assumption is a simplification, there is a vast empirical literature supporting the existence of partial insurance due to financial constraints Aiyagari (1994), Krusell and Smith (2006), and Guvenen (2011) to name a few studies.

[7]We assume that the capital of deceased households is used to pay an extra return equal to $1/\theta$ to those households who survive.

Finally, we denote by Φ the entire cross-sectional distribution of individual characteristics (a, ε, γ, β) and, together with the aggregate productivity shock z, summarize the aggregate state of the economy in each period.

2.2.2 Idiosyncratic countercyclical earning risk

The labor earning risk comes from two sources:

1. Idiosyncratic unemployment risk: in the spirit of Krusell and Smith (1998), the unemployment stochastic process follows a first-order Markov chain with transition matrices $\pi(\varepsilon, \varepsilon'|z, z')$. The matrices' dependence on the aggregate productivity transition allows the model to capture the effects of the business cycle on the persistence and incidence of unemployment.

2. Idiosyncratic efficiency risk: as it is common in the literature, the log-labor productivity of households follows a process with transitory and persistent component:[8]

$$\log(y_t) = \log(x_t) + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon).$$

   In line with Guvenen, Ozkan, and Song (2014), the persistent part follows an AR(1) process with persistence parameter $\phi \in [0, 1]$, and the innovations are drawn from a mixture of normal distributions whose parameters vary along with the business cycle:[9]

$$\log(x_t) = \phi \log(x_{t-1}) + \eta_t, \tag{1}$$

[8]This specification finds empirical support in Meghir and Pistaferri (2004), Storesletten, Telmer, and Yaron (2004b), Guvenen (2009), Meghir and Pistaferri (2011)

[9]The dependency between labor productivity and the business cycle is documented in Guvenen, Ozkan, and Song (2014), McKay (2017), Busch, Domeij, Guvenen, and Madera (2022), Busch and Ludwig (2021)

where

$$\eta_t = \begin{cases} \mathcal{N}(\mu_1(z_t), \sigma_1) & \text{with prob.} \quad p_1(z_t) \\ \mathcal{N}(\mu_2(z_t), \sigma_2) & \text{with prob.} \quad p_2(z_t) \\ \mathcal{N}(\mu_3(z_t), \sigma_3) & \text{with prob.} \quad p_3(z_t), \end{cases}$$

with $\sum_i p_i(z_t) = 1$, $p_i(z_t) \geq 0$, $z_t \in \mathcal{Z}$, and by normalization, $\mathbb{E}(\exp(\epsilon_t)) = 1$ and $\mathbb{E}(\exp(\eta_t)) = 1$. The latter normalization is because the primary interest of the paper is to analyze how fluctuations in the third moment of labor earning shocks change households saving behavior, while keeping constant the first moment. The log-labor efficiency process is discretized in n nodes $\mathcal{Y} = \{\gamma_1, \ldots, \gamma_n\}$. We assume that γ follows a Markov process with transition matrices $\pi(\gamma, \gamma'|z, z')$ which depends on the aggregate state of the economy. Due to the normalization described above, the discrete process for the idiosyncratic labor productivity shock satisfies $\sum_\gamma \Pi_z(\gamma)\gamma = 1$, for $z \in \mathcal{Z}$.

Both idiosyncratic shocks satisfy the law of large numbers. Consequently, only the aggregate shock z determines the share of households in each idiosyncratic state (ε, γ). These shares are denoted by $\Pi_z(\varepsilon)$ and $\Pi_z(\gamma)$, respectively.

2.2.3 Household decision problem

Given the distribution Φ and the aggregate shock z, a household with individual state variables $(a, \varepsilon, \gamma, \beta)$ solves the following recursive problem:

$$v(a,\varepsilon,\gamma,\beta;\Phi,z) = \max_{a'\geq 0,\, c\geq 0} \left\{ u(c) + \beta\theta \sum_{\{z',\varepsilon',\gamma'\}} \pi(\varepsilon',z'|\varepsilon,z)\pi(\gamma'|\gamma,z,z')\; v(a',\varepsilon',\gamma',\beta;\Phi',z') \right\}$$

$$s.t. \quad c + a' = \left[\frac{1-\delta+r(\Phi,z)}{\theta}\right] a + \big(1-\tau(z,\rho)\big) w(\Phi,z)\exp(\gamma)\varepsilon + b(\gamma;\Phi,z)(1-\varepsilon)$$

$$\Phi' = H(\Phi,z,z'),$$

where H represents the law of motion of the distribution of individual states. Notice that the prices $r(\Phi,z)$ and $w(\Phi,z)$, and unemployment insurance benefits $b(\gamma;\Phi,z)$ depends on the distribution of individual states and the aggregate shock.

2.2.4 Government and social security

The government implements a balanced budget unemployment insurance system:

$$\underbrace{\tau\left[\sum_{\gamma}\Pi_z(\gamma)\Pi_z(\varepsilon=1)w(\Phi,z)\exp(\gamma)\right]}_{\text{Tax revenue}} = \underbrace{\Pi_z(\varepsilon=0)\sum_{\gamma}\Pi_z(\gamma)b(\gamma;\Phi,z)}_{\text{Government spending}},$$

then the tax rate that balance the budget satisfies:

$$\tau(z,\rho) = \rho\left(\frac{\Pi_z(\varepsilon=0)}{1-\Pi_z(\varepsilon=0)}\right).$$

The tax rate depends on the business cycle because the aggregate productivity shock determines the unemployment rate.

2.3 Recursive competitive equilibrium

Given Φ, z and ρ, a recursive competitive equilibrium is characterized by a value function v, policy functions a' and c, pricing functions r and w, and an aggregate law of motion for Φ such that:

1. The value function v satisfies the Bellman equation. Also, given $r(\Phi, z)$ and $w(\Phi, z)$, a' and c are the associated policy functions.

2. Given $r(\Phi, z)$ and $w(\Phi, z)$, aggregate capital and labor satisfy:

$$\begin{aligned} r(\Phi, z) &= z f_K(K, L) \\ w(\Phi, z) &= z f_L(K, L) \end{aligned}$$

3. Markets clear for all (Φ, z):

$$\begin{aligned} L &= \Big(1 - \Pi_z(u)\Big) \sum_{\gamma \in \mathcal{Y}} \Pi_z(\gamma) \exp(\gamma) \\ K' &= \int a'(a, \varepsilon, \gamma, \beta; \Phi, z) \, \mathrm{d}\Phi(a, s, \gamma, \beta) \\ C &= \int c(a, \varepsilon, \gamma, \beta; \Phi, z) \, \mathrm{d}\Phi(a, \varepsilon, \gamma, \beta) \\ Y &= C + K' - (1 - \delta) K \end{aligned}$$

4. For all (Φ, z), the labor income tax rate τ is adjusted so that the Government follows a balanced budget policy.

5. The aggregate law of motion H is induced by the idiosyncratic exogenous stochastic and aggregate processes and by the optimal policy functions.

2.4 Computational aspects

In the recursive household decision problem, the cross-section distribution of individual characteristics Φ is an endogenous state variable. Households need to know how the distribution will evolve to forecast future prices. Unfortunately, the dimension of Φ is infinite, and numerical solutions to dynamic programming problems become more challenging as the number of state variables increases.

Thus, we solve the household problem using the Quasi-Aggregation algorithm proposed by Krusell and Smith (1998).[10,11] This algorithm assumes that agents are boundedly rational and perceive that current and future prices depend on a finite number of moments of the distribution of wealth. We assume that agents keep track only of the mean of the capital stock, allowing us to replace the aggregate law of motion for Φ with a log-linear law of motion for K that depends solely on the realization of z. Given the aggregate capital K and the aggregate shock z, a household with individual state $(a, \varepsilon, \gamma, \beta)$ solves the following recursive problem:

$$v(a,\varepsilon,\gamma,\beta;\Phi,z) = \max_{a'\geq 0,\, c\geq 0} \left\{ u(c) + \beta\theta \sum_{\{z',\varepsilon',\gamma'\}} \pi(\varepsilon',z'|\varepsilon,z)\pi(\gamma'|\gamma,z,z')\, v(a',\varepsilon',\gamma',\beta;\Phi',z') \right\}$$

$$\begin{aligned} s.t. \quad c + a' &= \left[\frac{1-\delta+r(\Phi,z)}{\theta}\right] a + \big(1-\tau(z,\rho)\big) w(\Phi,z)\exp(\gamma)\varepsilon + b(\gamma;\Phi,z)(1-\varepsilon) \\ \log(K') &= \psi_l + \kappa_l \log(K) \quad \text{if } z = z_l \\ \log(K') &= \psi_h + \kappa_h \log(K) \quad \text{if } z = z_h \end{aligned}$$

where ψ_l, ψ_h, κ_l and κ_h are constants to be determined using the Krusell and Smith (1998) method. We iterate on the Euler equation to solve the household decision problem, as in Maliar, Maliar, and Valli (2010).[12]

[10] To implement the Quasi-Aggregation algorithm, we simulate a continuum of agents using the method described in Ríos-Rull (1999). Simulating a continuum eliminates the sampling noise in some subgroups of households. See Algan, Allais, and Den Haan (2010), and Algan, Allais, Den Haan, and Rendahl (2014) for a discussion about the possible adverse effects of simulating a finite number of agents.

[11] See Appendix A.2 for details on the algorithm employed to simulate a continuum of agents.

[12] See Appendix A.3 for details of the Euler equation method.

3 Calibration

The model is calibrated to quarterly data. Table 1 reports the value, description, and source or target of the calibrated parameters.

Table 1: Calibration

Parameter	Value	Description	Source or Target
Basic Parameters			
σ	2	Coefficient of relative risk aversion	Standard value
$1-\theta$	0.5%	Probability of dying	Expected working lifetime: 50 years
δ	2.5%	Depreciation rate	Den Haan *et al.*, 2010
α	0.36	Capital share	Den Haan *et al.*, 2010
ρ	15%	Replacement rate	Den Haan *et al.*, 2010
Business cycle parameters			
(z_l, z_h)	$(0.9676,\ 1.0064)$	Aggregate productivity support	Krueger *et al.*, 2016a
$(\Pi_{z_l}(\varepsilon=0), \Pi_{z_h}(\varepsilon=0))$	$(8,39\%,\ 5,33\%)$	Unemployment rate	Krueger *et al.*, 2016a
$\pi(\varepsilon, \varepsilon' \vert z, z')$	See text	Transition matrix unemployment shock	Krueger *et al.*, 2016a
$\pi(z, z')$	See text	Transition matrix aggregate shock	Krueger *et al.*, 2016a
Discount factor parameters			
$\bar{\beta}$	0.9360	Mean discount factor	Capital to output ratio: 10.26
ν	0.0571	Discount factor dispersion	Wealth Gini coefficient: 0.78
n_β	7	Number of nodes for discretization	Carroll *et al.*, 2017
Idiosyncratic labor earnings shock parameters			
γ	See appendix	Idiosyncratic efficiency	Discretization
$\pi(\gamma' \vert \gamma, z, z')$	See appendix	Transition matrix of labor earnings process	Discretization

3.1 Parameters taken from literature

As is standard in the literature, we set the relative risk aversion parameter to $\sigma = 2$, the depreciation rate to $\delta = 2.5\%$, and the capital share to $\alpha = 0.36$. We set the probability of dying to $1 - \theta = 0.5\%$ for an expected working life of 50 years. Wet set the unemployment replacement rate to $\rho = 15\%$. To calibrate the parameters related to the business cycle, we follow Krueger, Mitman, and Perri (2016a), who defines a severe recession as one in which the unemployment rate exceeds 9% for at least one quarter. Its duration is determined by the number of quarters in which the unemployment rate exceeds 7%. Under this definition, over the period from 1948.I to 2014.III, the aggregate shock process reflects an average duration of 22 quarters for severe recessions.

The resulting transition matrix for the aggregate shock is:

$$\pi(z,z') = \begin{pmatrix} \rho_l & 1-\rho_l \\ 1-\rho_h & \rho_h \end{pmatrix} = \begin{pmatrix} 0.9545 & 0.0455 \\ 0.0090 & 0.9910 \end{pmatrix}$$

where ρ_l and ρ_h are the persistence parameters of severe recession and normal times, respectively. This parameterization implies that the invariant distribution for the aggregate technology shock is $\Pi(z) = [0.164, \;\; 0.836]$.

The idiosyncratic unemployment risk is determined by four employment-unemployment Markov transition matrices that depend on the economy aggregate state transition and are specified to reflect actual job search and separation rates in the CPS data. The unemployment transition matrices are taken directly from Krueger, Mitman, and Perri (2016a):

$$\pi(s,s'|z_l,z_l') = \begin{pmatrix} 0.3378 & 0.6622 \\ 0.0606 & 0.9394 \end{pmatrix}, \qquad \pi(s,s'|z_l,z_h') = \begin{pmatrix} 0.2220 & 0.7780 \\ 0.0378 & 0.9622 \end{pmatrix}$$

$$\pi(s,s'|z_h,z_l') = \begin{pmatrix} 0.3382 & 0.6618 \\ 0.0696 & 0.9304 \end{pmatrix}, \qquad \pi(s,s'|z_h,z_h') = \begin{pmatrix} 0.1890 & 0.8810 \\ 0.0457 & 0.9543 \end{pmatrix}$$

where the first element in each matrix corresponds to the probability that an unemployed household remains unemployed between the current period and the next.

3.2 Calibrated parameters and discretization

Following Krueger, Mitman, and Perri (2016a), the parameters that characterize the distribution of the discount factor $(\overline{\beta}, \nu)$ are calibrated to a Wealth Gini coefficient of 0.78 and a quarterly capital-to-output ratio K/Y of 10.26 (Carroll, Slacalek, Tokuka, and White, 2017). This targeted values require that $\bar{\beta} = 0.936$ and $\nu = 0.0571$. Thus, the discount factor is uniformly distributed between $[0.8789, \;\; 0.9931]$. The distribution is then discretized in 7 equidistant nodes as in Carroll, Slacalek, Tokuka, and White (2017).

The significant amount of heterogeneity in the discount factor deserves an explanation. Krusell and Smith (1998) have argued that just little heterogeneity in patience is sufficient to match the wealth distribution. Incorporating a small dispersion in the discount factor into their model increases the Gini coefficient of wealth from 0.25 to 0.82.

Why does Krusell and Smith (1998) require low heterogeneity in the discount factor to produce such an increment in wealth inequality? First, in Krusell and Smith (1998) the households are modestly risk-averse. Secondly, agents in their model face small aggregate productivity shocks, and the unemployment shock has low persistence, lasting two quarters on average (Hendricks, 2007). Consequently, there are no incentives to hold large amounts of precautionary savings. Then, a tiny dispersion in the discount factor is just enough to increase inequality in wealth holdings (Carroll, Slacalek, Tokuka, and White, 2017). In our model, agents have greater risk aversion and face more realistic labor earning risks (unemployment and efficiency shocks) and a greater dispersion and persistence in the aggregate productivity shock. Hence, small amounts of heterogeneity in the discount factor have almost no effect on the wealth distribution because households have substantial incentives to hold precautionary wealth. Therefore, we require a higher discount factor heterogeneity to generate a realistic wealth distribution featuring approximately 40% of agents with no or little wealth.

Finally, we discretize the log-labor idiosyncratic earning process into a first-order Markov chain modifying the method proposed in Civale, Díez-Catalán, and Fazilet (2016). The procedure provides the nodes $\{\gamma_1, \ldots, \gamma_n\}$, the invariant distributions $\Pi_z(\gamma)$ and the transition matrices $\pi(\gamma, \gamma'|z, z')$.[13]

4 Results

This section compares two versions of the model. As an extended part of the literature, the first version assumes that the labor efficiency process follows an AR(1) with innovations drawn from a normal distribution, approximating the original KMP

[13] The idiosyncratic labor earning process has an asymmetric distribution with high kurtosis. Therefore, its discretization is not trivial. Details of the procedure can be found in the appendix A.4.

model, but without the life cycle component and with different parameters values. We denote this version as the acyclical model. The second version assumes that the labor efficiency process follows an AR(1) with innovations drawn from a mixture of normal distributions so that the skewness of the distribution varies over the business cycle. Consequently, conditional on employment, significant earnings drops become more likely during recessions, whereas large upward movements become less likely. We denote this version as the cyclical model. The comparison of the models is in terms of (1) their ability to match the observed US cross-sectional wealth distribution and to what extent the models can reproduce the empirical joint distribution of income, consumption, and wealth. (2) The aggregate consumption, investment, and output response to negative technology shocks. We consider two types of technology shocks with different expected duration, a one-time negative technology shock and one negative aggregate shock as persistent as the Great Recession. Finally, (3) we compute individual welfare losses when the economy slips into severe economic downturns with an expected duration of 22 quarters and examine how these losses are distributed across households. Also, we measure the aggregate welfare losses of such types of severe recessions.

4.1 Income and consumption across the wealth distribution

The key elements that allow the KMP model to replicate the empirical wealth distribution are the discount factor heterogeneity and the idiosyncratic efficiency risk. The heterogeneity in preferences allows a non-negligible share of very patient households with a low propensity to consume to continue saving even at high levels of wealth. Similarly, it produces highly impatient households with little incentive to accumulate wealth, amplifying wealth inequality. Also, the inclusion of the stochastic and persistent labor efficiency process implies that households in the low productivity state remain in that state, on average, for a long time, thus making it difficult for them to accumulate wealth. In contrast, households in the high productivity state will accumulate wealth due to the fear of suffering a negative productivity shock. Considering those mentioned above, we explore to which extent the inclusion of countercyclical earnings risk, conditional on being employed, affects the ability of the model to replicate the observed US cross-sectional wealth distribution and the empirical joint distribution of disposable income, consumption, and wealth.

Table 2 reports key wealth distribution statistics computed from the data (2006 Panel Survey of Income Dynamics (PSID) and 2007 Survey of Consumer Finances (SCF)), the original KMP model, and our two model versions.[14] To make the comparison fair between models, all of them are calibrated to match the same Gini coefficient and capital-output ratio. The table shows several interesting facts. First, the wealth distribution is virtually identical in both of our models. Thus, the inclusion of countercyclical earning risk does not significantly alter the wealth distribution. Second, both of our models do a better job overall matching the empirical US wealth distribution. For instance, in the original KMP model, the middle class is too wealthy, and the wealthy are too poor compared to the data. However, the better fit comes at the cost of doing slightly worse at the bottom of the distribution. Finally, our models better match the wealth concentration at the top of the distribution. The data shows that the top 1% of wealthy households own 30% of overall wealth. In our two models, these households account for 23%, well above the 14% of the original KMP model.

Figure 1 presents the Lorenz curve for the wealth distributions of the data (SCF, 07) as well as the original KMP model and our two model versions. The figure clearly shows the patterns documented in the above paragraph. Our models do slightly worse matching the wealth distribution at the bottom, but much better fit at the top and very top of it.

Next, we test the ability of the cyclical model to replicate the joint distribution of disposable income, consumption expenditures, and wealth displayed in the PSID data.[15] Table 3 reports the share of disposable income and consumption expenditure by net

[14] SCF, PSID, and KMP model wealth distribution statistics are from Krueger, Mitman, and Perri (2016a).

[15] We do not report the acyclical model as their shares are nearly identical to those of the cyclical model.

Table 2: Wealth Distributions: Data v/s Models

% Share held by:	Data		Models		
	PSID, 06	SCF, 07	KMP	Acyclical	Cyclical
Q1	-0.9	-0.2	0.3	0.5	0.5
Q2	0.8	1.2	1.2	2.0	2.0
Q3	4.4	4.6	4.6	4.9	4.9
Q4	13.0	11.9	16.0	11.9	11.9
Q5	82.7	82.5	77.8	80.6	80.6
Top 10%	67.4	69.9	58.1	66.6	67.1
Top 5%	53.7	58.8	40.2	52.1	52.5
Top 1%	30.9	33.5	14.2	22.9	23.0
Gini	0.77	0.78	0.77	0.77	0.77

Notes: The table reports wealth distribution statistics computed from the data (2006 Panel Survey of Income Dynamics (PSID) and 2007 Survey of Consumer Finances (SCF)), the original KMP model, and our two model versions. SCF, PSID, and KMP model wealth distribution statistics are from Krueger, Mitman, and Perri (2016a).

Figure 1: Lorenz Curve: Data v/s Models

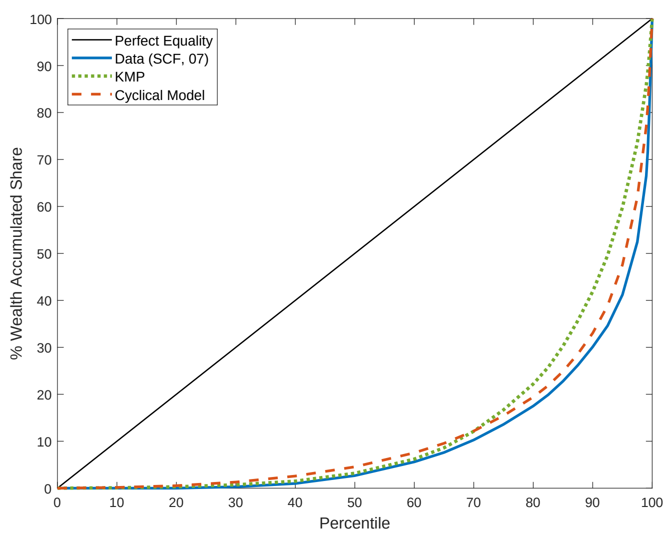

Notes: This figure displays the wealth distribution Lorenz Curve for SCF 07 data, the KMP model and the Cyclical Model.

worth computed from data (PSID), the original KMP model, and the cyclical model.[16] First, for all models in the table, there is a positive correlation between disposable income and consumption expenditure with net worth, as in the data. Second, in the data, the bottom two net worth quintiles account for 22.7% of overall consumption expenditures, while in the cyclical model, it is approximately 20%, improving upon the 17.9% of the original KMP model. This improvement is relevant because those quintiles have the most significant decline in consumption when the recession hits, which is essential for the macro response to aggregate shocks (Krueger, Mitman, and Perri, 2016a). The cyclical model generates wealth-poor households that are more consumption-rich than those of the original KMP model and wealth-rich households that are less consumption-rich, closing the discrepancy between the KMP model and the data.

Table 3: Selected Variables by Net Worth: Data v/s Models

	% Share of					
	Disposable Income			Expenditures		
Net Worth	PSID, 06	KMP	Cyclical	PSID, 06	KMP	Cyclical
Q1	8.7	6.0	8.3	11.3	6.6	7.4
Q2	11.2	10.5	13.1	12.4	11.3	12.4
Q3	16.7	16.6	18.2	16.8	16.6	17.7
Q4	22.1	24.3	24.4	22.4	23.6	24.4
Q5	41.2	42.7	36.1	37.2	42.0	38.1

Notes: This table reports the share of disposable income and consumption expenditure by net worth computed from data (PSID), the benchmark KMP model, and the two versions of our model. PSID and benchmark KMP model joint distribution statistics are from Krueger, Mitman, and Perri (2016a).

4.2 The dynamics of macroeconomic aggregates in severe recessions

The main finding of Krueger, Mitman, and Perri (2016a) is that an incomplete markets economy that generates a realistic wealth heterogeneity amplifies the aggregate consumption drop by a factor of two when a recession hits, relative to the representa-

[16] PSID and benchmark KMP model joint distribution statistics are from Krueger, Mitman, and Perri (2016a)

tive agent (RA) economy. The reason is that the KMP model, as observed in the data, generates a wealth distribution where nearly 40% of the population holds almost no wealth but represents an important part of aggregate consumption. When the aggregate economy slips into a recession, these wealth-poor households drastically reduce consumption for precautionary saving motives. In recessions, the probability of (transitory) earning loss due to an increase in the unemployment risk increases, and in the presence of incomplete markets, this increment in idiosyncratic risk generates a larger reduction in consumption relative to the RA economy.

The introduction of the fact that persistent large earnings drops become more likely than upward earnings movements during a recession into the KMP model should reinforce the precautionary saving motive, amplifying the decline in aggregate consumption and weakening its recovery (McKay, 2017; Amromin, De Nardi, and Schulze, 2018). In this section, we provide a quantitative answer to this conjecture.

Following Krueger, Mitman, and Perri (2016a), we consider two quantitative exercises. In both exercises, we take as an initial condition the wealth distribution produced after several realizations of normal times aggregate productivity shocks so that the distribution of wealth has stabilized. Then, the economy goes into a severe recession. The recession lasts only one quarter in the first exercise, returning to normal times afterward. In the second exercise, the economy goes into recession for one quarter. Then, it evolves stochastically according to the aggregate technology process transition matrix so that the recession will have an expected length of 22 quarters.[17] We simulate 10,000 aggregate productivity independent shock paths. Then for each period, we average across simulations the responses of the macroeconomic variables. We compute the impulse response functions of output, consumption, and investment for both experiments.

Figure 2 plots the impulse responses of aggregate consumption, investment, and output to a one-time recession type shock. The upper left panel displays the dynamics

[17]Note that this is different from the procedure carried out in Krueger, Mitman, and Perri (2016a), who simulate a recession that lasts 22 quarters.

of the technology shock, which is the same for both versions. The aggregate consumption drops by 3.5% in the cyclical model, while its decline is 3% in the acyclical model. The incorporation of countercyclical earnings risk, conditional on employment, amplifies the response of consumption by 0.5 percentage points. This magnitude is considerable. For example, Krueger, Mitman, and Perri (2016a) also find an additional drop in consumption of 0.5 percentage points in response to the same type of shocks when moving from the original low wealth inequality Krusell-Smith economy to their benchmark economy with realistic wealth distribution. The counterpart of the amplified consumption drop is a smaller investment decline in the cyclical model. Since output is used for consumption or investment, labor supply and efficiency are exogenous, and capital is a predetermined variable, the smaller decline in investment only translates into a slightly higher level of capital, generating virtually no difference in output dynamics between the two types of models under the one-period recession.

Figure 2: Impulse response to aggregate technology shock in two economies: One time technology shock

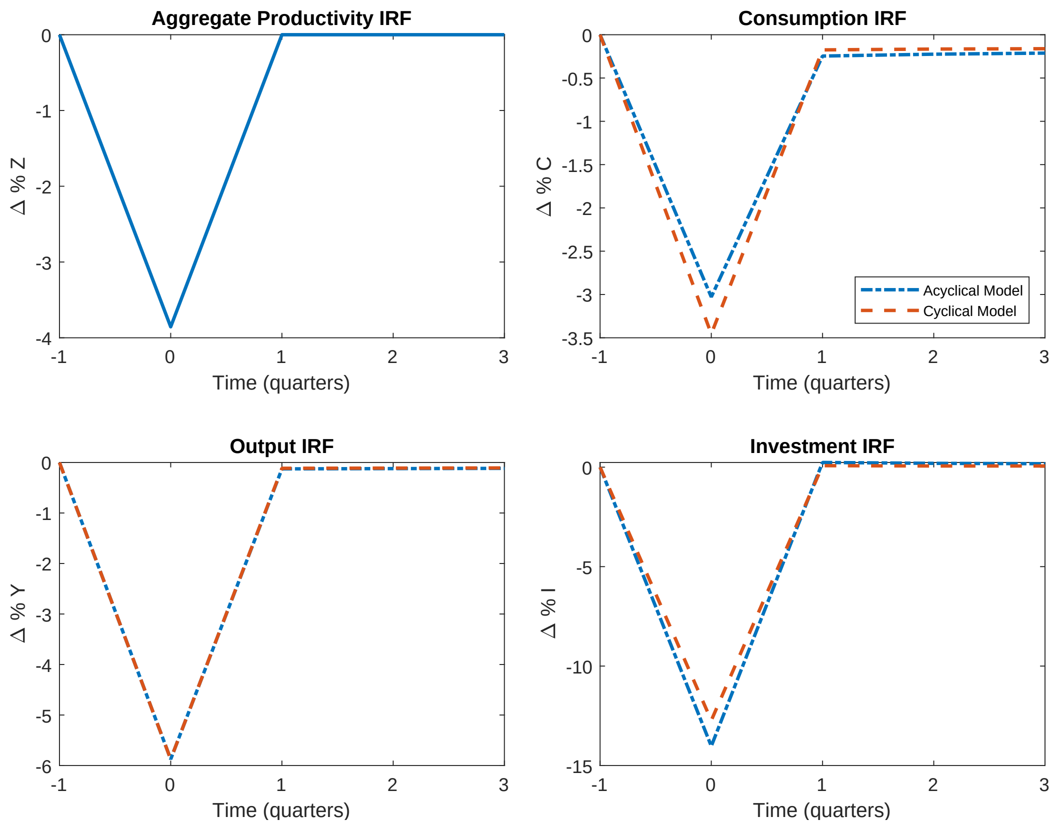

Notes: The figure displays dynamics of consumption, investment, and output in response to a one-time negative technology shock after a long sequence of normal times technology realizations for both versions of the model. The upper left panel displays the dynamics of the technology shock.

Figure 3 plots the average responses of the macroeconomic aggregates to a recession with an expected duration of 22 periods. The upper left panel shows the dynamics of the technology shock, which is the same for both versions of the model. The output dynamics for the two models are nearly identical; however, aggregate consumption and investment display different paths. Not only the magnitude of the drop in aggregate consumption differs, but also its dynamics. In the acyclical model, there is a smaller drop in aggregate consumption at the onset of the recession, but it continues to fall for several quarters. In the cyclical model, the drop in aggregate consumption is more prominent, and its growth after the onset of the recession has persistently languished. As of the tenth quarter, the dynamic of aggregate consumption is essentially the same for both types of models. The largest fall in aggregate investment for both economies occurs when the recession hits; nonetheless, the drop in investment is weaker in the cyclical economy as households increase their precautionary savings.[18]

[18] As in McKay (2017), we have assumed that the mean of the labor earnings shock distribution is constant over the business cycle. Therefore, the distribution median is larger in recessions than expansions to generate procyclical skewness. The above implies that fewer households draw negative income shocks in recessions than in expansions, which is economically counterintuitive. To address this concern in appendix A.7 we allow the mean of the idiosyncratic labor earnings shocks to vary over the business cycle, so more people draw negative income shocks in recessions than in expansions. Using this alternative process for the idiosyncratic earnings risk does not alter the main result: the inclusion of countercyclical earnings risk increases the reduction in consumption by 0.5 percentage points.

Figure 3: Impulse response to aggregate technology shock in two economies: Severe recession technology shock.

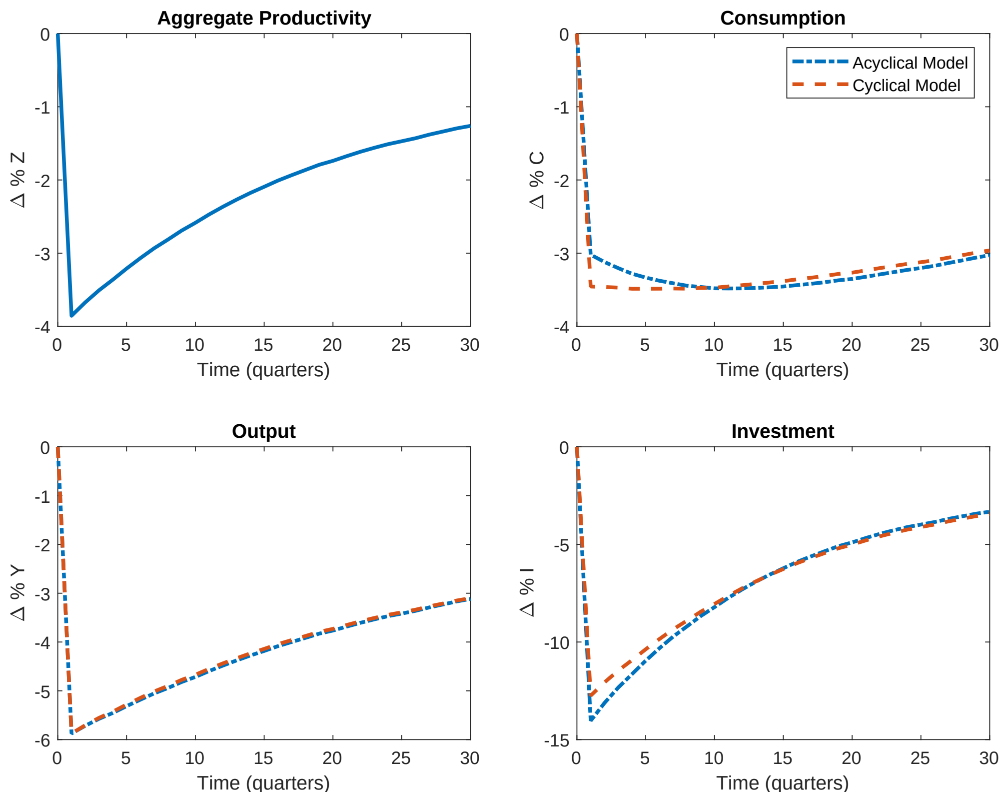

Notes: The figure displays the average responses of aggregate consumption, investment, and output to a recession that lasts on average 22 quarters for the two versions of the model. The upper left panel displays the dynamics of the technology shock.

What explains the different responses in aggregate consumption between the two economies? In the acyclical model, the only idiosyncratic risk that increases when the economy slips into a recession is the probability of unemployment, and its expected duration increases from 1.2 quarters in normal times to 1.5 quarters. The increased unemployment risk translates into a current and short-lived expected future income loss, which is easier to hedge. In contrast, in the cyclical model, in addition to unemployment risk, a long-lasting decline in earnings prospects increases during recessions. Because of the high persistence of the increased risk, households sharply cut consumption to form or increase their precautionary savings. In other words, what drives the difference in consumption dynamics is an increase in a highly persistent income risk because it is harder to ensure, not only for poor-wealth households but also for the rich-wealth.

To illustrate the latter mechanism, Table 4 shows the number of quarters that, on average, are needed to reach the 10th and 30th percentile of the invariant distribution of labor earnings, starting from the four lowest realizations of the labor efficiency shock. In recessions, a household in the lowest realization must wait on average 7.6 quarters to reach the 10th percentile of the invariant labor earnings distribution, while when the economy is in normal times, it must wait on average 6.1 quarters. In other words, the household must wait an average of 1.5 more quarters to reach the 10th percentile when the economy enters a recession. The difference is even more significant if a household wants to reach higher levels of labor earnings. For example, to reach the 30th percentile of the invariant labor earnings distribution, it must wait for an additional 1.9 quarters if the economy slips into a recession. Thus, the worst expected earnings prospects in recessions lead households to vigorously increase their precautionary savings, producing a more substantial decline in consumption in the economy with countercyclical earning risks.

Table 4: Number of quarters that, on average, are needed to reach the:

	First Income Decile		
Initial γ	Recession	Normal times	Δ
γ_1	7.6	6.1	1.5
γ_2	6.7	5.3	1.4
γ_3	6.0	4.6	1.4
γ_4	5.3	3.9	1.4
	Third Income Decile		
Initial γ	Recession	Normal times	Δ
γ_1	11.0	9.1	1.9
γ_2	10.2	8.3	1.9
γ_3	9.5	7.6	1.9
γ_4	8.7	6.9	1.8

Notes: The table shows the number of quarters that, on average, are needed to reach the 10th and 30th percentile of the invariant distribution of labor earnings, starting from the four lowest realizations of the labor productivity shock. We compute the expected quarters in recession and in normal times. The last column shows the difference in the number of quarters between the two aggregate states.

4.3 Welfare losses

One of the implications of the heterogeneity in consumption response across households is that welfare losses from large recessions will be unevenly distributed across the population. Considering that the bottom two wealth quintiles have almost no wealth but represent 23.7% of aggregate consumption, we should expect that welfare losses will be more pronounced in this group and, once aggregated across households, should be of a magnitude far greater than those found by Lucas (1987). As we know from the previous section, the inclusion of countercyclical earnings risk produces a more considerable drop in aggregate consumption in severe recessions, suggesting that welfare losses may be substantial. This section quantifies the welfare losses for our two model types. Below we explain how we measure individual and aggregate welfare losses.

In the spirit of Krueger, Mitman, and Perri (2016b) we measure household-specific welfare losses as the permanent percent increase in consumption that makes it indifferent between remaining in normal times z_h and experiencing a recession z_l with scaled-up consumption. This measure is known as the consumption compensating variation. Let $\lambda_{z_h,z_l}(a,\varepsilon,\gamma,\beta;K)$, be the required percentage consumption compensation for a household with individual characteristics $(a,\varepsilon,\gamma,\beta)$ to be willing to tolerate a severe recession today. Given a certain level of aggregate capital, the household-specific welfare losses when the economy transitions from normal times to severe recession are given by[19]

$$\lambda_{z_h,z_l}(a,\varepsilon,\gamma,\beta;K) = 100 \times \left[\left(\frac{v(a,\varepsilon,\gamma,\beta;K,z_l)}{v(a,\varepsilon,\gamma,\beta;K,z_h)} \right)^{\frac{1}{\sigma-1}} - 1 \right] \tag{2}$$

We measure aggregate welfare losses as the permanent increase in consumption that makes a household, under the veil of ignorance, indifferent between remaining in normal times z_h and experiencing a recession z_l with scaled-up consumption.[20] Let $\overline{\lambda}$ be the required percentage consumption compensation for an average household to be willing to tolerate a severe recession today. Given a certain level of capital K, the aggregate welfare losses when the economy transitions from normal times to a severe recession are given by[21]

$$\overline{\lambda} = 100 \times \left[\left(\frac{\displaystyle\int v(a,\varepsilon,\gamma,\beta;K,z')\, \mathrm{d}\Phi}{\displaystyle\int v(a,\varepsilon,\gamma,\beta;K,z)\, \mathrm{d}\Phi} \right)^{\dfrac{1}{\sigma-1}} - 1 \right] \tag{3}$$

Table 5 shows aggregate welfare losses and the fraction of households who experience welfare losses above certain thresholds (3% to 10%) for both model economies.

[19]See Appendix A.6.1 for the derivation of this result.

[20]Under the veil of ignorance, the household does not know how many assets it has nor what are its labor characteristics. It only knows how the distribution of individual states is after a long sequence of normal times aggregate shocks.

[21]See Appendix A.6.2 for the derivation of this result.

The aggregate welfare losses are around one percentage point greater in the cyclical model (4.1% vs. 3.1%). There is considerable heterogeneity in the welfare losses within both model versions. This heterogeneity is because these two models generate a wealth distribution that matches the empirically observed US wealth distribution. The wealth-rich households use their savings to hedge against the increased idiosyncratic risk that the recession brings, while households with little or no wealth increase their precautionary savings, cutting current and future consumption. Recall that in the previous section, we established that the inclusion of pro-cyclical earning risk, conditional on employment, exacerbates the consumption drop. Therefore, the substantial reduction in consumption generates bigger welfare losses and increases welfare heterogeneity as those poor-wealth households will be the ones that cut consumption the most.

Table 5: Household-specific and aggregate welfare losses from Great Recession

Model	Aggregate welfare loss (% of lifetime consumption)	% of households with loss $> 3\%$	$> 4\%$	$> 5\%$	$> 6\%$	$> 7\%$	$> 8\%$	$> 9\%$	$> 10\%$
1. Acyclical	3.11	58.45	29.24	11.15	3.40	1.09	0.72	0.51	0.34
2. Cyclical	4.10	60.21	37.28	23.28	15.72	10.97	7.69	5.62	3.85

Notes: The table shows aggregate welfare losses and the fraction of households who experience welfare losses above certain thresholds (3-10%) for both economies. Household-specific welfare losses are computed from equation (2), while aggregate welfare losses are computed from equation (3).

As shown in Figure 4, the cyclical model has a distribution of welfare losses with a right tail that is fatter and longer, which implies that a significant fraction of households suffers sizeable welfare losses. For example, around 23% of households experience losses bigger than 5% of lifetime consumption in the cyclical model, while it is about 11% in the acyclical model. Moreover, in the cyclical economy, around 4% of households suffer welfare losses greater than 10% of lifetime consumption, while only 0.35% suffer losses of this magnitude in the acyclical economy.

Figure 4: Distribution of Welfare Losses

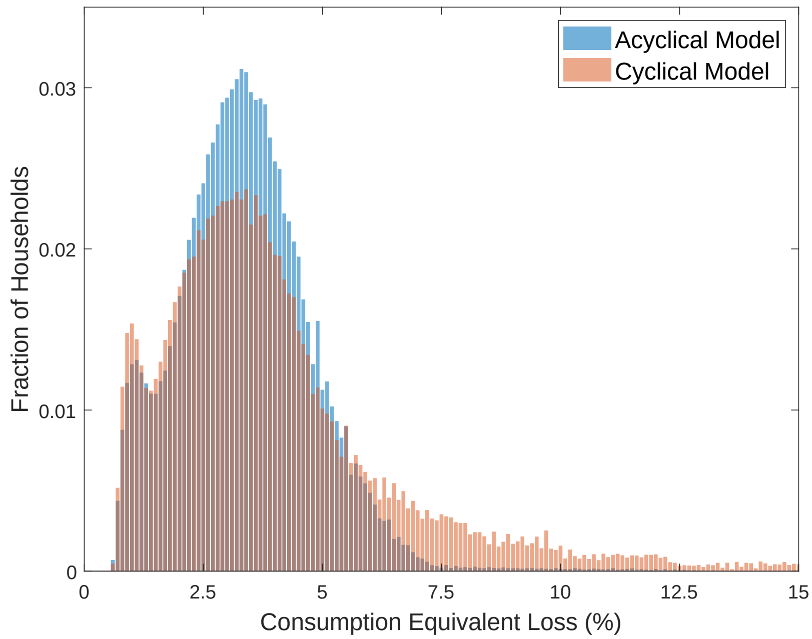

Notes: The figure shows household-specific welfare losses distribution for both versions of the model. Household-specific welfare losses are computed from equation (2).

Households that have been repeatedly unemployed and have been in low labor productivity states for multiple periods will not have enough opportunities to accumulate assets to protect themselves from the increased risk of a long-lasting earning decline. Less patient households, which have little wealth, will be seriously affected by the increased idiosyncratic risk since the short-term business cycle is the primary determinant of their lifetime utility Krueger, Mitman, and Perri (2016b). These types of households will be far away in the distribution's right tail.

Figure 5 plots the welfare losses from experiencing a Great Recession today against assets for currently employed households with the highest discount factor and numerous idiosyncratic efficiency levels for both economies.[22]

[22]Note that the figure plots welfare losses just for the change from z_h to z_l, without any idiosyncratic state change, such as the transition from employment to unemployment or high skill to low skill.

Figure 5: Welfare Losses from Great Recession by Asset Holdings: High β

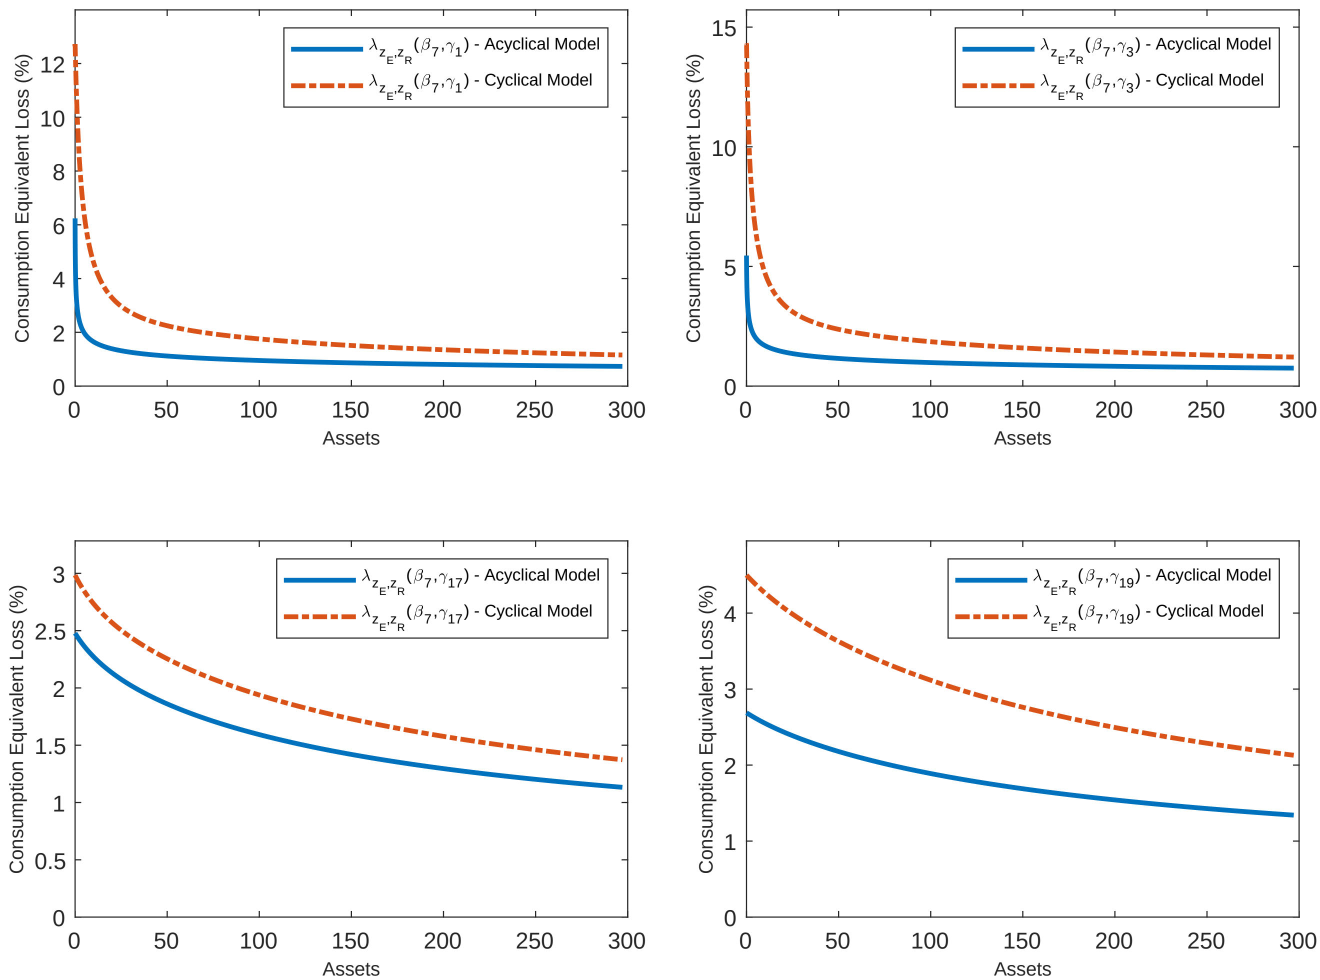

Notes: The figure plots the welfare losses from experiencing a Great Recession today against assets for currently employed households with the highest discount factor and numerous idiosyncratic efficiency levels for both economies. The top panel shows the results for low levels of efficiency while the bottom panel shows the results for high levels of efficiency.

The figure shows several interesting facts: (1) regardless of the model considered, the aggregate transition from normal times to severe recession, with the associated increase in idiosyncratic risks, is significantly costly in terms of welfare for all households, including the very rich. (2) The welfare loss is substantially large for those households with zero or little net worth and low idiosyncratic efficiency because losing one's job or experiencing a reduction in efficiency when holding little or no wealth implies a bigger consumption cut due to precautionary saving motives. (3) The inclusion of countercyclical earnings risk dramatically increases the welfare losses for the low-skilled households with little or no wealth. The difference in welfare losses between both models reduces as the household gets richer because, with more resources, households can hedge against the increased idiosyncratic risks. (4) Even for a moderate and high amount of wealth holdings, there is a noticeable difference in welfare losses between the models. Also, as efficiency increases, welfare losses rise too because those high productivity agents will cut consumption due to the fear of a long-lasting income decline during recessions.

5 Conclusion

This paper adds to a growing literature emphasizing the importance of countercyclical earning risks during recessions for consumption dynamics and welfare losses. We have argued that the inclusion of countercyclical labor income risk, conditional on employment, into a canonical real business cycle model with heterogeneous households and incomplete markets amplifies the response of aggregate consumption on impact by one percentage point to severe recessions such as the Great Recession of 2007-2009. Also, it significantly weakens the subsequent consumption growth. The worst labor earning prospects in recessions leads households to sharply cut consumption and increase their precautionary savings to insure themselves against the possibility of suffering a highly persistent fall in earnings during economic downturns.

We have also studied the welfare implications of countercyclical earning risks. The significant decline in aggregate consumption has its welfare counterpart. By reducing current and future consumption for precautionary motives, households experience sizeable welfare losses. Once aggregated, the welfare losses are about 4.1% of lifetime consumption. Furthermore, welfare losses vary enormously across household-specific characteristics. Those households with no or little wealth, representing approximately 40% of the population, experience higher losses as they cannot properly insure themselves against the increase in idiosyncratic earning risk during recessions.

In this work, the more consumption drops, the faster the recovery from recessions is. Moreover, this paper has no role for social insurance other than providing resources when unemployed. In reality, policymakers aim to stabilize output because of the endogenous feedback between consumption and economic activity. At least two straightforward extensions for future research could be taken to shed light on the importance of public policies employed during severe recessions. First, as Krueger, Mitman, and Perri (2016a) did, aggregate externality demands could model the negative loop between output and consumption. A drop in consumption would yield an additional output reduction, lowering aggregate wages, further exacerbating the consumption drop. Thus, social insurance programs aiming to reduce wealth inequality would stabilize consumption, decreasing the business cycle fluctuations. Second, in our model, households exogenously supply a unit of time to the labor market, though a proper calculation of the contribution of social insurance programs has to take into account the distortions generated by its financing via taxation. Incorporating endogenous labor supply choices into the model would give a more appropriate measure of the pros and cons of public policies.

References

Aiyagari, R. (1993). Explaining Financial Markets Facts: The Importance of Incomplete Markets and Transaction Costs. *Federal Reserve Bank of Minneapolis Quarterly Review* 17, 17-31.

Aiyagari, R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics* 109(3), 659-684.

Angelopoulus, K., Lazarakis S., and Malley, J. (2021). Cyclical Labour Income Risk in Great Britain. *Journal of Applied Econometrics*.

Algan, Y., Allais, O., and Den Haan, W. (2010). Solving the incomplete markets model with aggregate uncertainty using parameterized cross-sectional distributions. *Journal of Economic Dynamics and Control* 34(1), 59-68.

Algan, Y., Allais, O., Den Haan, W., and Rendahl, P. (2014). Solving and Simulating Models with Heterogeneous Agents and Aggregate Uncertainty. *Handbook of Computational Economics*, Volume 3, Chapter 6, pp. 277-324. Elsevier.

Amromin, G., De Nardi, M., and Schulze, K. (2018). Household Inequality and the Consumption Response to Aggregate Real Shocks. *Economic perspectives* 42(1).

Attanasio, O., and Pistaferri, L. (2016). Consumption Inequality. *Journal of Economic Perspectives* 30(2), 1-27.

Berger, D., and Vavra, J. (2015). Consumption Dynamics During Recessions. *Econometrica* 83(1), 101-154.

Busch C., Domeij D., Guvenen F., and Madera R. (2022). Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance. *American Economic Journal: Macroeconomics* 14(2), 207-242.

Busch, C., and Ludwig, A. (2021). Higher-Order Income Risk Over the Business Cycle. *Barcelona GSE Working Paper Series.* Working Paper 1159.

Caballero, R. (1990). Consumption Puzzles and Precautionary Savings. *Journal of Monetary Economics* 25, 113-136.

Cagetti, M. (2003). Wealth Accumulation Over the Life Cycle and Precautionary Savings. *Journal of Business and Economic Statistics* 21(3), 339-353.

Carroll, C. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis.*Quarterly Journal of Economics* 112(1), 1-55.

Carroll, C. (2000). Requiem for the Representative Consumer? Aggregate Implications of Microeconomic Consumption Behavior. *American Economic Review* 90(2), 110-115.

Carroll, C., Slacalek, J., Tokuka, K., and White, M. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics* 8(3), 977-1020.

Castañeda, A., Díaz-Giménez, J., and Ríos-Rull, J. (2003). Accounting for the U.S. Earnings and Wealth Inequality. *Journal of Political Economy* 111(4), 818-857.

Chatterjee, S., and Corbae, D. (2007). On the Aggregate Welfare Cost of Great Depression Unemployment. *Journal of Monetary Economics* 54, 1529-1544.

Civale, S., Díez-Catalán, L., and Fazilet, F. (2016). Discretizing a Process with Non-zero Skewness and High Kurtosis. *Working Paper, University of Minnesota*.

Cozzi, M. (2012). Risk Aversion Heterogeneity, Risky Jobs and Wealth Inequality. *Queen's Economics Department Workin Paper* 1286.

De Nardi, M. (2015). Quantitative Models of Wealth Inequality: A Survey. *NBER Working Paper Series*, Working Paper 21106.

De Nardi, M., French, E., and Benson, D. (2012). Consumption and the Great Recession. *Economics Perspectives* 36(1).

Deaton, A. (1991). Saving and Liquidity Constraints. *Econometrica* 59(5), 1221-1248.

Den Haan, W., Judd, K., and Juillard, M. (2010). Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty. *Journal of Economic Dynamics and Control* 34(1), 1-3.

Díaz-Giménez, J., and Prescott, E. (1992). Liquidity Constraints in Economies with Aggregate Fluctuations: A Quantitative Exploration. Research Department Staff Report 149, Federal Reserve Bank of Minneapolis.

Díaz-Giménez, J., Glover, A., and Ríos-Rull, J.-V. (2011). Facts on the Distributions of Earnings, Income, and Wealth in the United States: 2007 update. *Quarterly Review*.

Frederick, S., Loewenstein, G., and O'Donogue, T. (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature* 40, 351-401.

Garriga, C., and Hedlund, A. (2020). Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession. *American Economic Review*, 110(6), 1603-34.

Guvenen, F. (2007). Learning your earning: Are labor income shocks really very persistent?. *American Economic Review*, 97(3), 687-712.

Guvenen, F.(2009). An Empirical Investigation of Labor Income Processes. *Review of Economic Dynamics* 12(1), 55-79.

Guvenen, F. (2011). Macroecnomics with Heterogeneity: A Practical Guide. *Economic Quarterly* 97(3), 255-326.

Guvenen, F., Mckay, A., and Ryan, C. (2022). A Tractable Income Process for Business Cycle Analysis. *Working Paper*.

Guvenen, F., Ozkan, S., and Song, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy* 122(3), 621-660.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021). What Do Data on Millions of US Workers Reveal about Lifecycle Earnings Dynamics? *Econometrica* 89(5), 2303-39.

Heathcote, J., Storesletten, K. and Violante, G. (2009). Quantitative Macroeconomics with Heterogeneous Households. *Annual Review of Economics*, 1, 319-354.

Hendrickz, L. (2007). How Important is Discount Rate Heterogeneity for Wealth Inequality? *Journal of Economic Dynamics and Control* 31, 3042-3068.

Hugget, M. (1993). The Risk-Free Rate in Heterogeneous-Agent Incomplete-insurance Economies. *Journal of Economic Dynamics and Control* 17, 953-969.

Hugget, M. (1997). The One-Sector Growth Model with Idiosyncratic Shocks: Steady State and Dynamics. *Journal of Monetary Economics*, 39, 385-403.

Imrohoroglu, A. (2008). Welfare Costs of Business Cycles. *The New Palgrave Dictionary of Economics*, 2nd edition.

Japelli, T., and Pistaferri L. (2014). Fiscal Policy and MPC Heterogeneity. *American Economic Journal: Macroeconomics* 6(4), 107-136.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary Policy According to HANK. *American Economic Review*, 108(3), 697-743.

Krebs, T. (2007). Job displacement risk and the cost of business cycles. *American Economic Review*, 97(3), 664-686.

Krueger, D., Mitman, K., and Perri, F. (2016a). Macroeconomics and Household Heterogeneity. *Handbook of Macroeconomics*, Volume 2A, Chapter 11, pp. 843-921. Elsevier.

Krueger, D., Mitman, K., and Perri, F. (2016b). On the Distribution of the Welfare Losses of Large Recessions. *NBER Working Paper Series*, Working Paper 22458.

Krusell P., Mukoyama, T, Sahin, A., and Smith, A. (2009). Revisiting the welfare effects of eliminating business cycles. *Review of Economic Dynamics* 12(3), 393-404.

Krusell P., and Smith, A. (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy* 106(5), 867-896.

Krusell P. and Smith, A. (2006). Quantitative Macroeconomics Models with Heterogeneous Agents. *Advances in Economics and Econometrics: Theory and Applications*, Ninth World Congress, 2006.

Harmenberg., K. and Sievertsen, H. The Labor-Market Origins of Cyclical Income Risk. Manuscript

Lawrance, E. (1991). Poverty and the Rate of Time Preference: Evidence from Panel Data. *Journal of Political Economy* 99, 54-77.

Low, H., Meghir, C., and Pistaferri, L. (2010). Wage risk and employment risk over the life cycle. *American Economic Review*, 100(4), 1432-67.

Lucas, R. Jr. (1987). Models of Business Cycles. *Yrjo Jahnsson Lecture Series, Helsinki, Finland.* Blackwell.

Maliar, L., Maliar, S., and Valli, F. (2010). Solving the Incomplete Markets Model with Aggregate Uncertainty using the Krusell-Smith Algorithm. *Journal of Economic Dynamics and Control* 34(1), 42-49.

McKay, A. (2017). Time-varying Idiosyncratic Risk and Aggregate Consumption Dynamics. *Journal of Monetary Economics* 88(1), 1-14.

McKay, A., and Reis, R. (2021). Optimal Automatic Stabilizers. *Review of Economic Studies*, 1-32.

Meeuwis, M. (2021). Idiosyncratic Income Risk, Precautionary Savings, and Asset Prices. *Working Paper.*

Meghir, C., and Pistaferri, L. (2011). Earnings, Consumption and Life Cycle Choices. *Handbook of Labor Economics*, Volume 4B, Chapter 9, pp. 773-854. Elsevier.

Meghir, C., and Pistaferri, L. (2004). Income Variance Dynamics and Heterogeneity. *Econometrica* 72(1), 1-32.

Nakajima, M., and Smirnyagin, V. (2019). Cyclical Labor Income Risk. *Working Paper*.

Petev, I., Pistaferri, L., and Saporta, I. (2012). Consumption and the Great Recession: An Analysis of Trends, Perceptions, and Distributional Effects. *Analyses of the Great Recession*.

Pistaferri, L. (2016). Why Has Consumption Remained Moderate after the Great Recession? *Working Paper*.

Rawls, J. (1999). A Theory of Justice. *Belknap Press: An Imprint of Harvard University Press*, 2nd edition.

Ríos-Rull, J. (1999). Computation of Equilibria in Heterogenous Agent Economies. *Computational Methods for the Study of Dynamic Economies*, Chapter 11, pp. 238-264. Oxford University Press.

Salgado, S., Guvenen, F., and Bloom, N. (2019). Skewed Business Cycles. *NBER Working Paper Series*, Working Paper 26565.

Scheinkman, J., and Weiss, L. (1986). Borrowing Constraints and Aggregate Economic Activity. *Econometrica* 54(1), 23-45.

Storesletten, K., Telmer, C., and Yaron, A. (2001). The Welfare Cost of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk. *European Economic Review* 45(1), 1311-1339.

Storesletten, K., Telmer, C., and Yaron, A. (2004a). Consumption and Risk Sharing. *Journal of Monetary Economics* 51(1), 609-633.

Storesletten, K., Telmer, C., and Yaron, A. (2004b). Cyclical Dynamics in Idiosyncratic Labor Market Risk. *Journal of Political Economy* 112(3), 695-717.

Zeldes, S. (1989a). Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence. Quarterly Journal of Economics 104(1), 275-298.

Zeldes, S. (1989b). Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy* 97(1), 305-246

A Appendix

A.1 Estimation of Earnings Process

Guvenen, Ozkan, and Song (2014) uses data on earnings histories from 1978 to 2011 from the US Social Security Administration records to estimate the process of log-labor productivity. The estimated process features parameters that change over the business cycle. Given that our quantitative model is at a quarterly frequency, we convert the annual process estimated in Guvenen, Ozkan, and Song (2014) to a quarterly process that, once aggregated to a yearly basis, minimizes the distance from selected moments. Those moments try to capture how the distribution of income changes varies over the business cycles, in specific, how the tails of the distribution change while the median varies little. We target the difference between the 10th and 90th percentile, the difference between the 50th and 10th percentile, the difference between the 90th and 50th percentile and the Kelley skewness of 1, 3, and 5-year income changes, distinguishing between periods of expansion and contraction and the persistence of the process. The quarterly process is the following:

$$\log(y_t) = \phi \log(y_{t-1}) + \eta_t,$$

where η_t follows a mixture of normal distributions:

$$\eta_t = \begin{cases} \mathcal{N}(\mu_1(z_t), \sigma_1) & \text{with prob.} \quad p_1(z_t) \\ \mathcal{N}(\mu_2(z_t), \sigma_2) & \text{with prob.} \quad p_2(z_t) \\ \mathcal{N}(\mu_3(z_t), \sigma_3) & \text{with prob.} \quad p_3(z_t), \end{cases}$$

Note that the means and probabilities of the mixture change along with the economy's state. As in McKay (2017), McKay and Reis (2021), and Guvenen, Mckay, and Ryan (2022) we normalize $\mathbb{E}(\exp(\eta_t)) = 1$ and $\mathbb{E}(\exp(\epsilon_t)) = 1$, so changes in the skewness of the distribution will have no effects on the first moment.

The procedure to determine the values of the parameters of the quarterly process is the following:

1. We simulate a long time series of aggregate shocks using the matrix outlined in section 3.2. Then, using the annual process, we simulate a large panel of earning histories. For expansions and contractions, compute the 10th, 50th, and 90th percentile of 1, 3, and 5-year changes and the Kelley skewness.
2. We use a global optimization procedure to find the quarterly process's parameter values that, once aggregated at annual frequency, minimize the percentage difference between the moments generated by the annual process with those obtained by aggregation of the quarterly process. To search the parameter values, we use the Particle Swarm Optimization algorithm.[23]

The procedure give us the following parameters for the quarterly process.

Table A.1: Estimated parameters for the process at quarterly frequency

ρ	$p_{1,R}$	$p_{2,R}$	$p_{3,R}$	$p_{1,E}$	$p_{2,E}$	$p_{3,E}$	σ_1	σ_2	σ_3	$\mu_{1,R}$	$\mu_{2,R}$	$\mu_{3,R}$	$\mu_{1,E}$	$\mu_{2,E}$	$\mu_{3,E}$
0.986	0.863	0.072	0.065	0.782	0.133	0.043	0.161	0.189	0.328	0.021	-0.457	0.086	-0.0361	-0.075	0.327

[23] https://la.mathworks.com/help/gads/particle-swarm-optimization-algorithm.html

A.1.1 Goodness of fit

We denote Δ_i as the i-th difference, LXY as the difference between the Xth and Yth percentile, and $\mathcal{K}$ the Kelley skewness.

Table A.2: Annual and quarterly aggregated moments percentage difference: $\Delta_1 \log(y)$

Recession				Expansion			
$L9010$	$L5010$	$L9050$	$\mathcal{K}$	$L9010$	$L5010$	$L9050$	$\mathcal{K}$
0.069	0.073	0.0645	0.055	0.072	0.069	0.074	0.040

Table A.3: Annual and quarterly aggregated moments percentage difference: $\Delta_3 \log(y)$

Recession				Expansion			
$L9010$	$L5010$	$L9050$	$\mathcal{K}$	$L9010$	$L5010$	$L9050$	$\mathcal{K}$
-0.029	-0.032	-0.026	-0.025	-0.018	-0.015	-0.021	-0.038

Table A.4: Annual and quarterly aggregated moments percentage difference: $\Delta_5 \log(y)$

Recession				Expansion			
$L9010$	$L5010$	$L9050$	$\mathcal{K}$	$L9010$	$L5010$	$L9050$	$\mathcal{K}$
-0.048	-0.050	-0.047	-0.019	-0.038	-0.037	-0.039	-0.011

Figure A.1: Density of $\Delta_1 \log(y)$: quarterly aggregated estimation

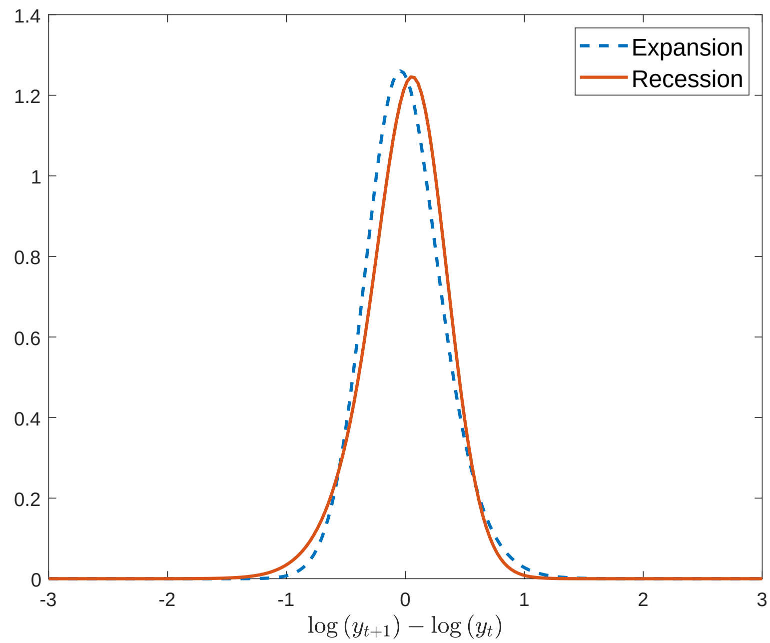

Figure A.2: Density of $\Delta_3 \log(y)$: quarterly aggregated estimation

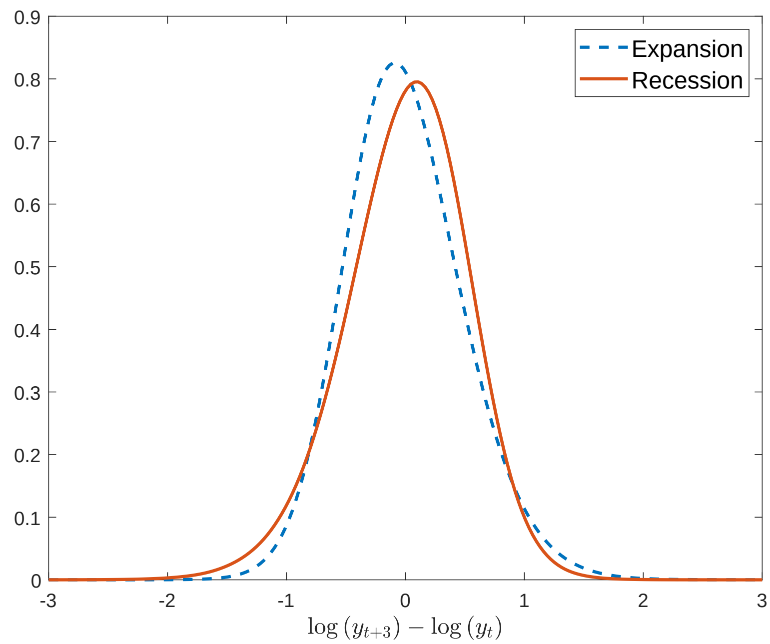

Figure A.3: Density of $\Delta_5 \log(y)$: quarterly aggregated estimation

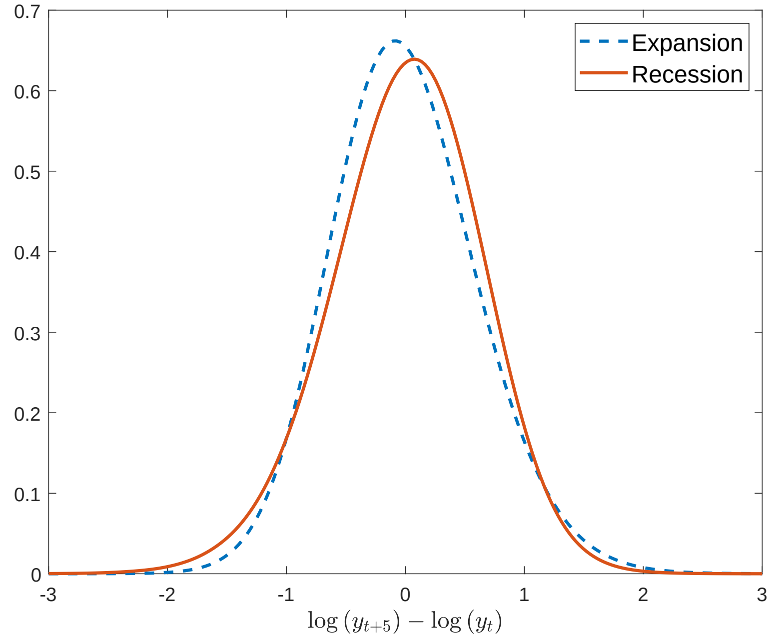

A.2 Simulation of a Continuum of Agents

In this section, we describe the procedure of Ríos-Rull (1999) and then adapted it by Algan, Allais, Den Haan, and Rendahl (2014) to simulate a continuum of agents. In this procedure, the CDF is approximated with a linear spline, meaning that a uniform distribution between grid points is assumed. At each node κ, we calculate the capital stock at the beginning of the period x, which would lead to the value of κ. That is, x is the inverse of κ according to the asset policy function. The algorithm proceeds as follows:

1. Grid: construct a grid and define the capital distribution at the beginning of period $t = 0$ as follows:
   (a) $\kappa_0 = 0$ and κ_i, for $i = 1, \ldots, I$.
   (b) Let $p_{\omega,0,t}$ be the share of agents in state $\omega \in \Omega = \{0,1\} \times \{\gamma_1, \ldots, \gamma_n\} \times \{\beta_1, \ldots, \beta_m\}$ that have capital stock equal to zero at the beginning of the period t.
   (c) For $i > 0$, let $p_{\omega,i,t}$ be the mass of agents with a capital stock greater than κ_{i-1} and less than κ_i. It is assumed that this mass of individuals is uniformly distributed over points on the grid.
   (d) Note that:

$$\sum_{i=0}^{I} p_{\omega,i,t} = 1.$$

   Denote this initial distribution by $P_{\omega,t}(k)$.

2. Distribution at the end of the period: calculate the level of assets such that the agent chooses a capital equal to κ_i for the next period. Denote this level by $x_{\omega,i,t}$. By definition:

$$a'(x_{\omega,i,t}, \varepsilon, \gamma, \beta; K_t, z_t) = \kappa_i$$

   For any point on the grid, the cumulative density function for the end of period t

for agents with state ω is given by:

$$F_{\omega,i,t} = \int_0^{x_{\omega,i,t}} dP_{\omega,t}(k) = \sum_{i=0}^{\bar{i}_{\omega,t}} p_{\omega,i,t} + \frac{x_{\omega,i,t} - \kappa_{\omega,\bar{i}_{\omega,t}}}{\kappa_{\bar{i}_{\omega,t}+1} - \kappa_{\bar{i}_{\omega,t}}} p_{\omega,\bar{i}_{\omega,t}+1,t}$$

where $\bar{i}_{\omega,t} = \bar{i}(x_{\omega,i,t})$ is the largest value of i such that $\kappa_i \leq x_{\omega,i,t}$. The second equality follows from the assumption that $P_{\omega,t}$ is uniformly distributed over points on the grid.

3. Initial distribution in the next period: let $g_{\omega_t,\omega_{t+1},z_t,z_{t+1}}$ be the mass of agents with state ω_t today and with state ω_{t+1} next period, conditional on the values of z_t, z_{t+1}. So, for each combination of z_t and z_{t+1}, it follows that:

$$\sum_{\omega_t \in \Omega,\ \omega_{t+1} \in \Omega} g_{\omega_t,\omega_{t+1},z_t,z_{t+1}} = 1$$

From this, we get

$$P_{\omega,i,t+1} = \sum_{\omega_t \in \Omega} \left(\frac{g_{\omega_t,\omega_{t+1}}}{\sum_{\omega_t \in \Omega} g_{\omega_t,\omega_{t+1}}} \right) F_{\omega,i,t}$$

and

$$p_{\omega,0,t+1} = P_{\omega,0,t+1}$$
$$p_{\omega,i,t+1} = P_{\omega,i,t+1} - P_{\omega,i-1,t+1}$$

A.3 Iterating on the Euler Equation

We iterate on the Euler equation proposed in Maliar, Maliar, and Valli (2010) to obtain the policy functions. This method has the advantage of being faster to compute and is more accurate than value function iteration. One drawback, however, is that its convergence is less stable, so it should be used with a damping parameter, as we will show.

The Euler equation, the budget constraint, the borrowing constraint, and the Kuhn-Tucker are, respectively:

$$
\begin{aligned}
&c^{-\sigma} + h = \beta \mathbb{E}\left[c'^{-\sigma}(1-\delta+r')\right] && \text{Euler equation} \\
&c + a' = \left[\frac{1-\delta+r}{\theta}\right] a + (1-\tau) w \exp(\gamma)\varepsilon + b(1-\varepsilon) && \text{Budget constraint} \\
&a' \geq 0 && \text{Borrowing constraint} \\
&h \geq 0,\ ha' = 0 && \text{Kuhn-Tucker conditions.}
\end{aligned}
$$

Form the budget constraint:

$$c(a',\varepsilon,\gamma,\beta) = \left[\frac{1-\delta+r}{\theta}\right] a + (1-\tau) w \exp(\gamma)\varepsilon + b(1-\varepsilon) - a'$$

Guessing a' and computing $a'' = a'(a')$, we get an expression to iterate on:

$$c(\tilde{a}',\varepsilon,\gamma,\beta)^{-\sigma} = h + \beta \mathbb{E}\left[c(a'',\varepsilon',\gamma',\beta)^{-\sigma}(1-\delta+r')\right] \tag{4}$$

$$
\begin{aligned}
\Leftrightarrow \quad \tilde{a}' = &\left[\frac{1-\delta+r}{\theta}\right] a + (1-\tau) w \exp(\gamma)\varepsilon + b(1-\varepsilon) \\
&- \left\{ h + \beta \mathbb{E}\left[\frac{1-\delta+r'}{\left(\left[\frac{1-\delta+r'}{\theta}\right] a' + (1-\tau') w' \exp(\gamma')\varepsilon' + b(1-\varepsilon') - a''\right)^{\sigma}}\right]\right\}^{-\frac{1}{\sigma}}
\end{aligned}
\tag{5}
$$

where $h \equiv h(a,\varepsilon,\gamma,\beta;K,z)$, $a' \equiv a'(a,\varepsilon,\gamma,\beta;K,z)$ and $a'(a') \equiv a'(a'(a,\varepsilon,\gamma,\beta;K,z))$.

Formally, the solution algorithm is as follows:

1. Choose the relevant space for asset holdings $a \in [0, a_{\max}]$ and for aggregate capital $K \in [K_{\min}, K_{\max}]$, then discretize these intervals to generate the grids. Given that the asset policy function has more curvature near the borrowing constraint but is almost linear in high levels of wealth, we placed more grid points at low asset holdings using the following formula outlined in Maliar, Maliar, and Valli (2010):

$$a_j = \left(\frac{j}{J}\right)^{\vartheta} a_{\max}, \qquad \text{for } \; j = 0, 1, \ldots, J$$

where $J + 1$ is the number of grid points, and ϑ controls the concentration of points in the beginning. As ϑ increases, more grid points are placed at the beginning, and fewer are placed towards the end of the grid. In practice, we use $\vartheta = 8$. We use an evenly spaced grid for aggregate capital because the asset policy function is almost linear in that dimension.

2. Guess an initial policy function for capital $a'(a, \varepsilon, \gamma, \beta; K, z)$ for the values on the grid.

3. For each point in the grid $(a, \varepsilon, \gamma, \beta; K, z)$, plug the policy function $a'(a, \varepsilon, \gamma, \beta; K, z)$ on the right side of equation (5), set the Lagrangian multiplier to equal zero, and compute the new policy function for capital, $\tilde{a}'(a, \varepsilon, \gamma, \beta; K, z)$. For any point in the grid such that $\tilde{a}'(a, \varepsilon, \gamma, \beta; K, z)$ is not in the range $[0, a_{max}]$, set $\tilde{a}'(a, \varepsilon, \gamma, \beta; K, z)$ equal to the value of the corresponding limit.

4. Update the policy function using the following formula:

$$\tilde{\tilde{a}}'(a, \varepsilon, \gamma, \beta; K, z) = (1 - \omega)\tilde{a}'(a, \varepsilon, \gamma, \beta; K, z) + \omega a'(a, \varepsilon, \gamma, \beta; K, z)$$

where $\omega \in (0, 1]$ is a damping parameter. We use a small value for ω so the new guess for the policy is less prone to oscillations, which could hinder the convergence.

5. Iterate steps 2-4 until convergence:

$$\left|\left|\tilde{\tilde{a}}'(a,\varepsilon,\gamma,\beta;K,z) - a'(a,\varepsilon,\gamma,\beta;K,z)\right|\right|_{\max} < 10^{-7}$$

A.4 Discretizing an Asymmetric Process with High Kurtosis

The idiosyncratic labor earnings process has asymmetric distribution and high kurtosis, consequently, its discretization is not trivial. We discretize this process using a first order markov chain through a modification of the method proposed in Civale, Díez-Catalán, and Fazilet (2016). Specifically, consider T realizations of the aggregate shock $Z = \{z_1, \ldots, z_T\}$ and $T \times M$ realizations of the idiosyncratic shock $Y_m = \{y_{1,m}, \ldots, y_{T,m}\}$, $m = 1, \ldots, M$.

The set of nodes of the calibrated markov chain is denoted by $\Gamma = \{\gamma_1, \ldots, \gamma_N\}$. This set of nodes is the same for both recessions and expansions. The nodes are chosen to match specific moments of Y. Denote by $M(Y,Z)$ the vector of moments of the original process and by $\hat{M}(\Gamma, Z)$ the same vector of moments but generated by the discrete process.

The procedure is

(i) Choose the number of nodes N.

(ii) Choose node values, $\Gamma = \{\gamma_1, \ldots, \gamma_N\}$.

(iii) For chosen nodes, map the realizations Y into a sequence of discrete realizations $\{x_1, \ldots, x_T\}$ as follows:

$$x_t = \underset{\gamma \in \Gamma}{\text{argmin}} \quad |y_t - \gamma|, \quad t = 1, \ldots, T.$$

(iii) Given $X = \{x_1, \ldots, x_T\}$ get $\hat{M}(\Gamma, Z)$ and compute:

$$F = \left[\frac{M(Y,Z) - \hat{M}(\Gamma,Z)}{M(Y,Z)}\right]' W \left[\frac{M(Y,Z) - \hat{M}(\Gamma,Z)}{M(Y,Z)}\right]$$

where W is a weight matrix. We use identity.

(iv) Find $\Gamma = \{\gamma_1, \ldots, \gamma_N\}$ to minimize F iterating steps (ii) and (iii). To find the nodes we use Simulated Annealing.

Denote by $\pi(\gamma, \gamma') = \{\pi(\gamma, \gamma'|z_l, z_l),\ \pi(\gamma, \gamma'|z_h, z_l),\ \pi(\gamma, \gamma'|z_l, z_h),\ \pi(\gamma, \gamma'|z_h, z_h)\}$ the transition matrices for labor earning process for the different combinations of transitions between aggregate states. Given the nodes $\Gamma = \{\gamma_1, \ldots, \gamma_N\}$ we can obtain the transition matrices.

(i) For transitions between aggregate states $(z_t = z_l, z_{t+1} = z_l)$ and $(z_t = z_h, z_{t+1} = z_h)$ compute the transition probability between idiosyncratic state i to j as:

$$\pi(\gamma_i, \gamma_j | z_t, z_{t+1}) = \frac{\sum_{t=1}^{N-1} \mathbb{1}\left(x_t = \gamma_i, x_{t+1} = \gamma_j, z_t = z, z_{t+1} = z'\right)}{\sum_{t=1}^{N-1} \mathbb{1}\left(x_t = \gamma_i, z_t = z, z_{t+1} = z'\right)}$$

wherer $\mathbb{1}(\cdot)$ is the indicator function.

(ii) To compute the probabilities of the transitions matrices associated with expansion to recession, and vice versa, we follow closely Krusell and Smith (1998), extended to 19 possible realizations of the idiosyncratic earnings shock.

Finally, to ensure consistency, the probabilities are adjusted so that the following equations hold $\forall (z, z') \in \mathcal{Z} \times \mathcal{Z}$:

$$\begin{pmatrix} \Pi_{z'}(\gamma_1) \\ \Pi_{z'}(\gamma_2) \\ \vdots \\ \Pi_{z'}(\gamma_{N-1}) \\ \Pi_{z'}(\gamma_N) \end{pmatrix} = \begin{pmatrix} \pi(\gamma_1, \gamma_1|z, z') & \pi(\gamma_2, \gamma_1|z, z') & \ldots & \pi(\gamma_{N-1}, \gamma_1|z, z') & \pi(\gamma_N, \gamma_1|z, z') \\ \pi(\gamma_1, \gamma_2|z, z') & \pi(\gamma_2, \gamma_2|z, z') & \ldots & \pi(\gamma_{N-1}, \gamma_2|z, z') & \pi(\gamma_N, \gamma_1|z, z') \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \pi(\gamma_1, \gamma_{N-1}|z, z') & \pi(\gamma_2, \gamma_{N-1}|z, z') & \ldots & \pi(\gamma_{N-1}, \gamma_{N-1}|z, z') & \pi(\gamma_N, \gamma_{N-1}|z, z') \\ \pi(\gamma_1, \gamma_N|z, z') & \pi(\gamma_2, \gamma_N|z, z') & \ldots & \pi(\gamma_{N-1}, \gamma_N|z, z') & \pi(\gamma_N, \gamma_N|z, z') \end{pmatrix} \begin{pmatrix} \Pi_z(\gamma_1) \\ \Pi_z(\gamma_2) \\ \vdots \\ \Pi_z(\gamma_{N-1}) \\ \Pi_z(\gamma_N) \end{pmatrix}$$

$$\sum_{j=1}^{N} \pi(\gamma_i, \gamma_j | z, z') = 1, \quad \forall\, i = 1, \ldots, N$$

A.5 Recovering the Value Function

The model is solved via iteration of the Euler equation, however, to analyze welfare the value function is needed. This section explains how to retrieve the value function from the policy function.

1. Obtain $a'(a, \varepsilon, \gamma, \beta; K, z)$, $c(a, \varepsilon, \gamma, \beta; K, z)$ by iterating the Euler equation.
2. Guess $v^0(a, \varepsilon, \gamma, \beta; K, z)$ and compute $v^i(a, \varepsilon, \gamma, \beta; K, z)$, $i = 0, 1, 2, \ldots$ using the following equation:

   $$v^{i+1}(a, \varepsilon, \gamma, \beta; K, z) = u\Big(c(a, \varepsilon, \gamma, \beta; K, z)\Big) + \beta\theta\mathbb{E}\left[v^i\left(a'(a, \varepsilon, \gamma, \beta; K, z), \varepsilon', \gamma', \beta; K', z'\right)\right]$$

   We use an interpolation with splines to evaluate the value function in points outside the grids individual assets and aggregate capital.
3. Repeat step 2 until $||v^j - v^{j-1}||_{\max} < Tol$.

A.6 Individual and Aggregate Welfare Losses

A.6.1 Individual welfare losses quantification

Consider the lifetime utility of a household with individual characteristics $(a, \varepsilon, \gamma, \beta)$ that follows the optimal policy under the aggregate state (K, z_h):

$$v(a, \varepsilon, \gamma, \beta; K, z_h) = \mathbb{E}_0\left[\sum_{t=0}^{\infty}(\beta\theta)^t \frac{c_t^{1-\sigma}}{1-\sigma}\right]$$

Next, consider the previous household, but the aggregate state of the economy has changed to z_l and the household is compensated by scaling up its consumption by a factor λ in every t and at every node of the event tree. Its lifetime utility is given by

$$
\begin{aligned}
v(a,\varepsilon,\gamma,\beta;K,z_l,\lambda) &= \mathbb{E}_0\left[\sum_{t=0}^{\infty}(\beta\theta)^t\frac{\Big((1+\lambda)c_t\Big)^{1-\sigma}}{1-\sigma}\right] \\
&= (1+\lambda)^{1-\sigma}\mathbb{E}_0\left[\sum_{t=0}^{\infty}(\beta\theta)^t\frac{c_t^{1-\sigma}}{1-\sigma}\right] \\
&= (1+\lambda)^{1-\sigma}v(a,\varepsilon,\gamma,\beta;K,z_l)
\end{aligned}
$$

For the household to be indifferent between normal times or the economy entering a severe recession but receiving compensation, we must find the value of λ such that

$$
\begin{aligned}
v\left(a,\varepsilon,\gamma,\beta;K,z_l,\lambda\right) &= v(a,\varepsilon,\gamma,\beta;K,z_h) \\
\Leftrightarrow \quad (1+\lambda)^{1-\sigma}v(a,\varepsilon,\gamma,\beta;K,z_l) &= v(a,\varepsilon,\gamma,\beta;K,z_h)
\end{aligned}
$$

Therefore, the scaling factor λ, as a percentage is

$$
\lambda_{z_h,z_l}(a,\varepsilon,\gamma,\beta) = 100\times\left[\left(\frac{v(a,\varepsilon,\gamma,\beta;K,z_l)}{v(a,\varepsilon,\gamma,\beta;K,z_h)}\right)^{\frac{1}{\sigma-1}}-1\right] > 0,
$$

as long as $v(a,\varepsilon,\gamma,\beta;K,z_l)/v(a,\varepsilon,\gamma,\beta;K,z_h) < 1$, which is true under a severe recession. In other words, if $\lambda_{z_h,z_l}(a,\varepsilon,\gamma,\beta) > 0$ and $\sigma > 1$, the household gets a positive compensation.

A.6.2 Aggregate welfare losses quantification

The average welfare in the economy with aggregate capital K and state of aggregate shock z_h is given by

$$
\int v(a,\varepsilon,\gamma,\beta;K,z_h)\,\mathrm{d}\Phi = \int \mathbb{E}_0\left[\sum_{t=0}^{\infty}(\beta\theta)^t\frac{c_t^{1-\sigma}}{1-\sigma}\right]\,\mathrm{d}\Phi
$$

Next, consider the previous economy, but the aggregate state of the economy has changed to z_l where all individuals are compensated by scaling up its consumption by a factor $\overline{\lambda}$ in every t and at every note of the event tree. Its lifetime utility is given by

$$\int v\left(a, \varepsilon, \gamma, \beta; K, z_l, \overline{\lambda}\right) \ \mathrm{d}\Phi = \int \mathbb{E}_0 \left[\sum_{t=0}^{\infty} (\beta\theta)^t \frac{\left(\left(1+\overline{\lambda}\right) c_t \right)^{1-\sigma}}{1-\sigma} \right] \ \mathrm{d}\Phi$$

Under the Veil of Ignorance, by how much would each agent in the economy have to be compensated, in terms of equivalent consumption units, to be indifferent between normal times or severe recession getting a compensation $\overline{\lambda}$? We must find the value of $\overline{\lambda}$ such that

$$\begin{aligned} \int v\left(a, \varepsilon, \gamma, \beta; K, z_l, \overline{\lambda}\right) \ \mathrm{d}\Phi &= \int v(a, \varepsilon, \gamma, \beta; K, z_h) \ \mathrm{d}\Phi \\ \Leftrightarrow \quad (1+\overline{\lambda})^{1-\sigma} \int v(a, \varepsilon, \gamma, \beta; K, z_l) \ \mathrm{d}\Phi &= \int v(a, \varepsilon, \gamma, \beta; K, z_h) \ \mathrm{d}\Phi \end{aligned}$$

Therefore, $\overline{\lambda}$, as a percentage is

$$\overline{\lambda} = 100 \times \left[\left(\frac{\int v(a, \varepsilon, \gamma, \beta; K, z_l) \ \mathrm{d}\Phi}{\int v(a, \varepsilon, \gamma, \beta; K, z_h) \ \mathrm{d}\Phi} \right)^{\frac{1}{\sigma-1}} - 1 \right]$$

A.7 Alternative idiosyncratic earnings risk process

The innovations of the persistent component are drawn from a mixture of normal distributions whose parameters vary with the business cycle:

$$\eta_t = \begin{cases} \mathcal{N}(\mu_1(z_t), \sigma_1) & \text{with probability} \quad p_1(z_t) \\ \mathcal{N}(\mu_2(z_t), \sigma_2) & \text{with probability} \quad p_2(z_t) \end{cases}$$

with $\sum_i p_i(z_t) = 1$, $p_i(z_t) \geq 0$, and $z_t \in \mathcal{Z} = \{z_l, z_h\}$.

There are two important features to notice. First, the process does not impose restrictions on the mean of the innovations of the persistent component. Therefore, process's mean and median is larger in expansions than in recessions. This approach may seem flawed at first sight, but it finds support in the empirical literature. For instance, Guvenen, Ozkan, and Song (2014) argued that the cyclical nature of labor earning shocks arises from the behavior of the tails of its distribution, which oscillate back and forth along the business cycle, displaying, therefore, procyclical skewness. Since the median exhibits small movements, the tail swings are the main driver of the changes in the mean of labor income shocks. Thus, recessions are best described as a modest negative shock to the median and a large negative shock to the skewness of the distribution of idiosyncratic labor income shocks, with little changes in its variance (Guvenen, Ozkan, and Song, 2014).

Second, due to the procyclical nature of the skewness of the distribution of the idiosyncratic earning shocks, if we impose some restriction on its mean, we will be assuming that more households receive more modestly, albeit very persistent, positive shocks in recessions than in expansions, which is economically counterintuitive. To illustrate this comment, consider the permanent component of the idiosyncratic efficiency process of Meeuwis, M. (2021), which follows a similar specification as McKay

(2017),

$$
\begin{aligned}
\log(x_t) &= \log(x_{t-1}) + \eta_t, \\
\text{where} \quad \eta_t &\sim \begin{cases} \mathcal{N}(\mu_{1,t}, \sigma_1) & \text{with probability} \quad p_1 \\ \mathcal{N}(\mu_{2,t}, \sigma_2) & \text{with probability} \quad p_2 \\ \mathcal{N}(\mu_{3,t}, \sigma_3) & \text{with probability} \quad 1 - p_1 - p_2 \end{cases} \\
\text{and} \quad \mu_{1,t} &= \overline{\mu}_t, \\
\mu_{2,t} &= \overline{\mu}_t + \mu_2 - x_t, \\
\mu_{3,t} &= \overline{\mu}_t + \mu_3 - x_t,
\end{aligned}
$$

where $\mu_2 < 0 < \mu_3$ and x_t is a risk factor that shifts the tails of the distribution of earnings growth. The term $\overline{\mu}_t$ is such that $\mathbb{E}[\exp(\eta_t)] = 1, \ \forall t$. This seemingly innocuous normalization implies that in recessions, where the term x_t grows, the distribution of η_t has a larger median than in expansions, where the term x_t decreases. Consequently, more people draw positive shocks in recessions than in expansions.

Meeuwis, M. (2021) presents the logarithm of the distribution density of η_t to argue that the shifts of the tails in recessions and expansions produce a small change in the median. However, a closer look at the distribution density of η_t reveals the opposite, as Figure A.4 shows.

A.7.1 A One-Time Negative Technology Shock

To aid the comparison in the one-time negative technology shock, we add the response of an representative agent economy. Figure A.5 plots aggregate consumption, investment, and output impulse responses to a one-time recession shock. The upper left panel displays the dynamics of the technology shock, which drops further in the representative agent and acyclical model to match the same initial output drop in recessions that generates the cyclical model. The figure reveals that the one-time shock induces a consumption drop of 2.97% in the cyclical model, 2.43% in the acyclical model, and 1.96% in the representative agent economy. Thus, the same output decline gener-

Figure A.4: Density of annual earnings change in Meeuwis, M. (2021).

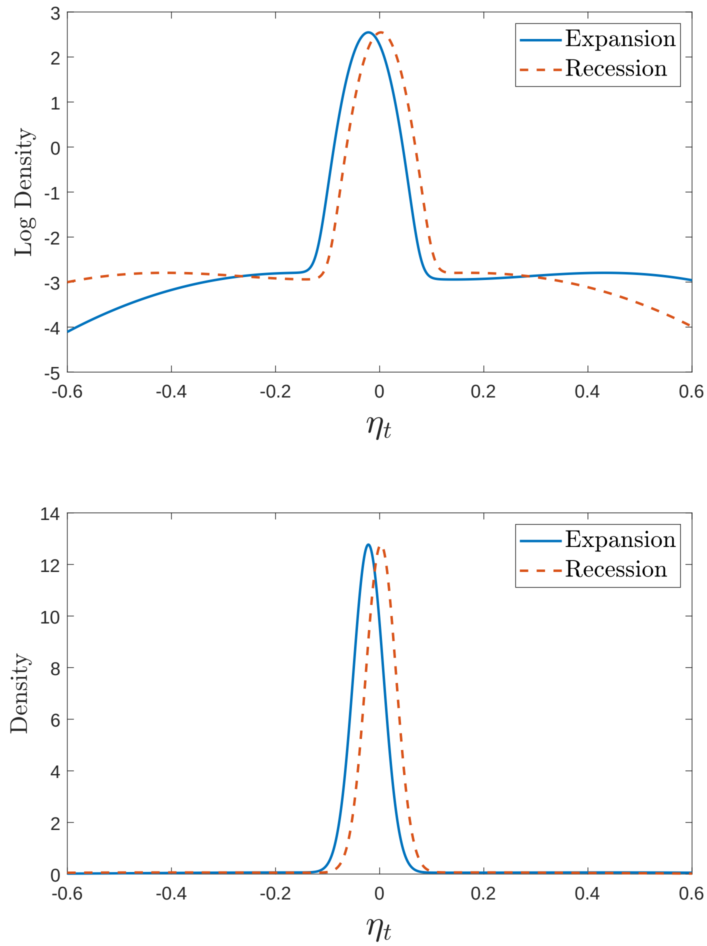

Note: The figure displays the distribution density of η_t.
Source: own simulation using Meeuwis, M. (2021) process and estimated parameters.

ates a consumption drop 0.54 percentage points larger (or 22% larger) in the cyclical model than in the acyclical model. Also, the acyclical model generates a consumption drop 0.47 percentage points larger (or 24% larger) than the representative agent model. Thus, conditional on employment, cyclical labor earning risk is as relevant as modeling economies that produce realistic wealth inequality for accounting for the sharp consumption drop observed in the data. Moreover, since the output is used for consumption or investment, and labor supply and efficiency are exogenous, there is a smaller fall in investment in the cyclical model relative to its acyclical counterpart. This smaller fall in investment translates into a slightly higher level of capital, generating virtually no difference in output dynamics between the acyclical and cyclical models in the one-period recession experiment.

Figure A.5: Impulse Response: one-time negative technology shock

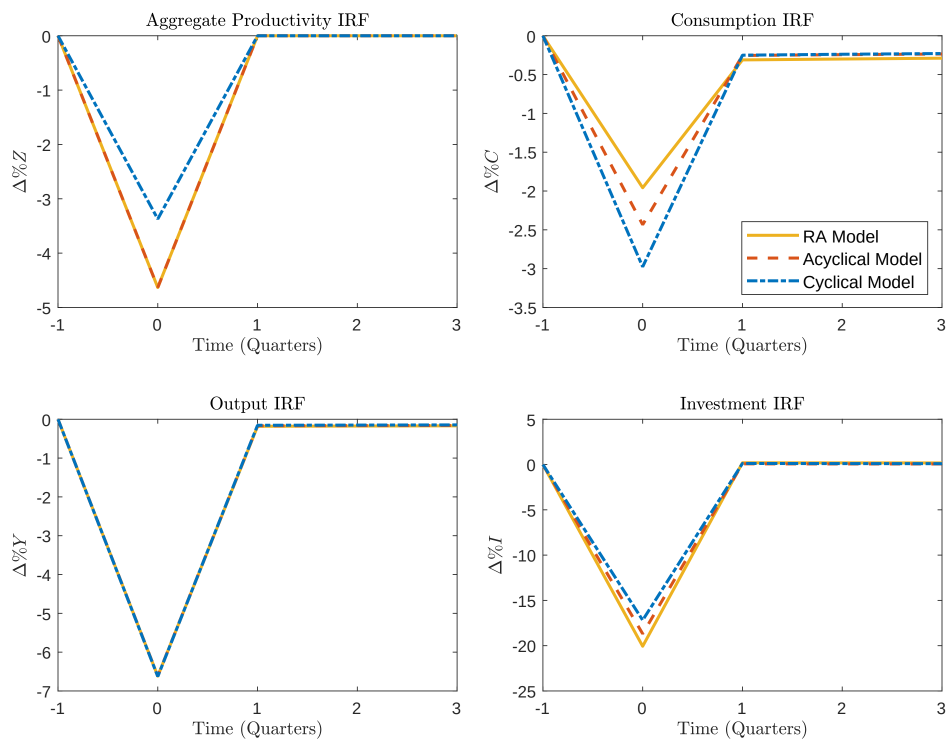

Notes: The figure displays dynamics of consumption, investment, and output in response to a one-time technology shock after a long sequence of normal times technology realizations for both versions of the model. The upper left panel displays the dynamics of the technology shock.

A.7.2 Expected Severe Recession-Type Shock

Figure A.6 plots the average responses of the macroeconomic aggregates to a recession with an expected duration of 22 periods. The upper left panel shows the dynamics of the technology shock, which drops further in the acyclical model to match the same initial output drop in the cyclical model when the economy slips into a recession. The output dynamics for the two models are nearly identical; however, aggregate consumption and investment display different paths. Not only the magnitude of the drop in aggregate consumption differs, but also its dynamics. In the acyclical model, there is a smaller drop in aggregate consumption at the onset of the recession, and it continues to fall for several quarters. In the cyclical model, the drop in aggregate consumption is more profound and continues to fall but not as strongly as in the acyclical model. As of the twenty-second quarter, the dynamic of aggregate consumption is essentially the same for both types of models. The largest fall in aggregate investment for both economies occurs when the recession hits. Nonetheless, the drop in investment is weaker in the cyclical economy as households increase their precautionary savings relative to the acyclical Model.

Figure A.6: Impulse Response: Severe recession technology shock.

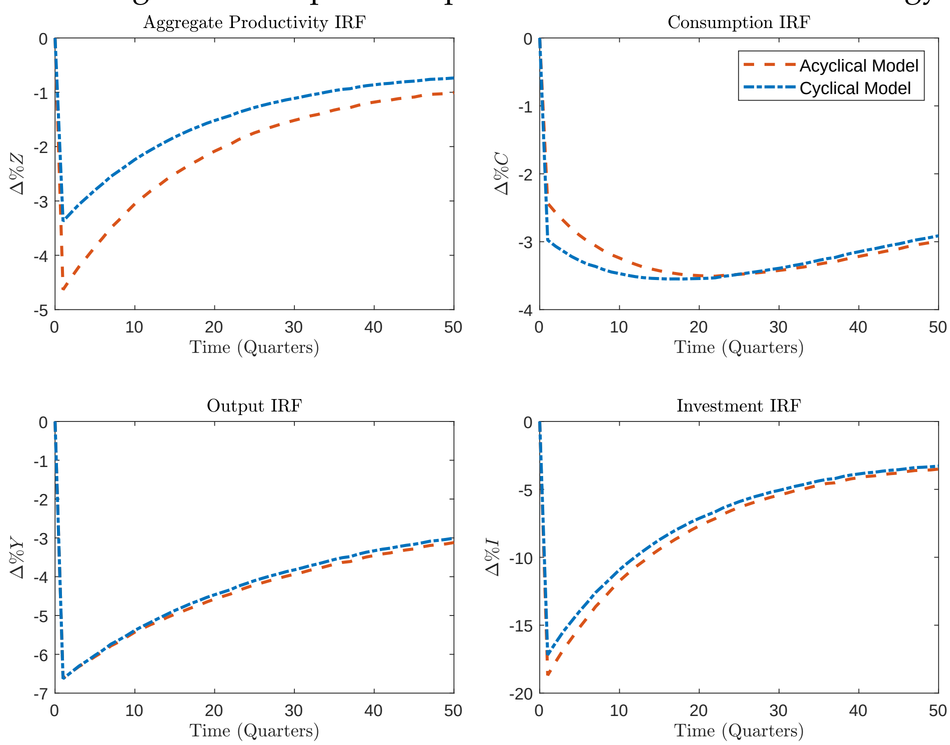

Notes: The figure displays dynamics of consumption, investment, and output in response to a one-time technology shock after a long sequence of normal times technology realizations for both versions of the model. The upper left panel displays the dynamics of the technology shock.

What explains the different responses in aggregate consumption between the two economies? In the acyclical model, only the probability of unemployment increases when the economy slips into a recession, and its expected duration increases from 1.2 quarters in normal times to 1.5 quarters in recessions. The increased unemployment risk translates into a current and short-lived expected future income loss, which is easier to hedge. In contrast, in the cyclical model, there is an increase in long-lasting decline in earnings prospects during recessions in addition to unemployment risk. Because of the high persistence of the increased risk, households cut consumption sharply to increase their precautionary savings. In other words, the difference in consumption dynamics reflects an increase in a highly persistent income risk that is more difficult to insure against, not only for poor-wealth households but also for the wealthiest, as we have shown.