Cyclical Labor Income Risk
and Consumption Dynamics in Large Recessions

TESIS PARA OPTAR AL GRADO DE
MAGÍSTER EN ECONOMÍA

Alumno: Francisco Díaz-Valdés
Profesor Guía: Eduardo Engel

Santiago, Mayo de 2023

Cyclical Labor Income Risk and Consumption Dynamics in Large Recessions*

Francisco Díaz-Valdés†

Universidad de Chile

May 29, 2023

Abstract

Recent empirical studies on labor income dynamics depict recessions as times when employed households face a higher risk of large and long-lasting income declines. Motivated by this evidence, we build a real business cycle model with heterogeneous agents and incomplete markets, in which aggregate productivity fluctuates, and idiosyncratic labor income risk varies along the business cycle. We use the model to investigate how much higher labor income risk during the Great Recession can account for the observed sharp and prolonged drop in US aggregate consumption. Compared to a model with just unemployment risk, we find that including cyclical labor earning risk amplifies the initial response of aggregate consumption to severe recessions by one percentage point (from 2% to 3%), and its subsequent recovery is significantly weakened. Also, we corroborate that the result holds even if TFP remains constant. However, we argue that TFP fluctuations should be considered, as we show that they contribute substantially to the sharp drop in consumption and its sluggish recovery.

*I thank Álvaro Aguirre, Eduardo Engel, and Álvaro García for their invaluable guidance and suggestions. I also thank Carolina Allende, Bernardo Candia, and Patricio Romero for their helpful comments. I gratefully acknowledge financial support from ANID - Subdirección de Capital Humano/Magíster Nacional/2018 - 22182052.

†fdiazvalde@fen.uchile.cl

1 Introduction

The Great Recession of 2007-2009 was the worst recession the United States had experienced since World War II. The decline in economic activity was profound and affected all macroeconomic aggregates, particularly consumption, which exhibited intriguing dynamics. Specifically, consumption during the Great Recession fell sharply, and its recovery languished (Pistaferri, 2016). Without considering the fall in housing and asset prices, how much does the business cycle with incomplete markets research program explain the initial sharp drop in consumption and its subsequent weak recovery?

The influential paper by Krusell and Smith (1998) showed that real business cycle models with incomplete markets featuring low wealth inequality produce consumption dynamics that differ little from their complete market counterparts. Consequently, the consumption drop in these models will be barely larger than their complete market counterparts in response to negative real macroeconomic shocks, such as total factor productivity (TFP) shocks. However, Krueger, Mitman, and Perri (2016a) (henceforth KMP) found that incomplete market models could produce large consumption drops in response to negative TFP shocks if they feature a realistic wealth distribution where approximately 40% of households hold little or no wealth. The intuition is that households with low wealth have a higher marginal propensity to consume, so they strongly reduce consumption for precautionary saving motives when the economy falls into recession. Nevertheless, while providing valuable insights, the KMP model still cannot reproduce the deep consumption drop and its protracted recovery, as they abstract from important changes in the labor earnings distribution during the Great Recession (Amromin, De Nardi, and Schulze, 2018).

For instance, the latest literature on labor income dynamics describes recessions as times when employed households face greater risks to their labor income.[1] These income risks arise from both short-lived and long-lived sources. On the one hand, short-term

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

risk takes the form of unemployment risk, the incidence and duration of which increases during recessions. On the other hand, although currently not fully understood, long-term risks appear to originate from changes in idiosyncratic wages, hours worked, or both. For example, Guvenen, Ozkan, and Song (2014) and Busch, Domeij, Guvenen, and Madera (2022) have documented that in recessions, individuals with strong labor market attachment are more likely to experience large and long-lasting negative labor income shocks, while the likelihood of experiencing large and long-lasting upward income shocks decreases.[2] In other words, the distribution of these highly persistent labor income shocks exhibits procyclical skewness, meaning that in recessions, the left tail enlarges while the right tail shrinks, yet the median varies slightly relative to either tail.

In this paper, we seek to answer three unexplored interrelated questions: (i) to what extent do household wealth inequality and idiosyncratic cyclical labor income risk amplify the initial aggregate consumption response to large negative TFP shocks? (ii) how much do these cyclical risks exacerbate the initial consumption response across the wealth distribution? Furthermore, (iii) after the initial consumption drop, to what extent do these cyclical long-lasting income risks account for the slow recovery of consumption?

To answer these questions, we extend the work of Krueger, Mitman, and Perri (2016a) and McKay (2017) by building a canonical real business cycle model that features aggregate productivity fluctuations, incomplete markets, and cyclical idiosyncratic income risks. We also revise the current procedure to incorporate idiosyncratic cyclical risk to labor income into the standard incomplete markets model.

As mentioned above, the cyclicality of labor income risk originates from two sources in the model. First, in the spirit of Krusell and Smith (1998), unemployment is exogenously determined, and its duration and persistence are higher during recessions. However, unemployment is typically considered a transitory shock to labor income (McKay, 2017). Second, conditional on employment, the persistent labor income risk distribu-

[2]By strong labor market attachment, we mean individuals who do not go through extended periods of unemployment. For instance, Guvenen, Ozkan, and Song (2014) focuses on prime-age male workers earning more than a time-varying threshold, and Busch, Domeij, Guvenen, and Madera (2022) found substantial cyclicality in the earnings growth distribution even for continuously employed workers.

tion exhibits procyclical skewness. Because understanding the procyclical nature of these long-lasting shocks is still in its early stages, we follow McKay (2017) and McKay and Reis (2021) by forgoing attempts to unfold the source from which it originates. Rather, we take as given the procyclical skewness of labor income shocks and then assess the implications for the response of aggregate consumption and consumption across the wealth distribution.

We contribute to the literature by revisiting and modifying the most common procedure used to incorporate procyclical income risk into standard incomplete market models and by investigating to what extent the interplay between aggregate shocks, persistent time-varying idiosyncratic risk, and household wealth inequality can account for the observed consumption dynamics during the Great Recession. Also, we study the consumption response across wealth quintiles.

To carry out our research, we proceed in three steps. First, we explore if our model can match the share of net worth held by US wealth distribution quintiles. A wealth distribution that resembles its empirical counterpart is necessary because low-wealth households possess higher propensities to consume from their income. Therefore, as expected incomes decline in recessions, they will reduce consumption sharply for precautionary motives. Our model is consistent with the empirical fact that the share of wealth held by the two poorest quintiles is virtually zero.

Second, we evaluate the impulse responses of macroeconomic aggregates to severe economic downturns, focusing mainly on the response of aggregate consumption. On this subject, we differentiate from the work of McKay (2017) in three aspects. (i) In his model, the mean of the idiosyncratic labor earnings shock is constant over the business cycle. Consequently, if the distribution of idiosyncratic earnings shock displays procyclical skewness, its median needs to be larger in recessions than in expansions. Therefore, fewer households draw negative shocks in recessions than in expansions. Adopting an alternative approach, we do not keep the mean of the distribution constant, implying that more people draw negative income shocks in recessions than in expansions. (ii) He holds TFP constant in his main result. Thus, his model generates small output declines, implying minor changes in factor prices. Our analysis considers negative TFP shocks and their role in lowering household income by reducing factor prices, which could be an important mechanism to consider as it contributes to lower current and future disposable income, decreasing consumption further. (iii) Finally, he does not study the consumption response throughout the distribution of wealth. Studying the response across the wealth distribution could provide new insights into understanding the consumption dynamics when the economy slips into a severe recession.

We evaluate the response of macroeconomic aggregates by performing two experiments: (i) a one-time negative technology shock and (ii) a stochastic duration-type shock that lasts 22 quarters on average. The first experiment gauges the initial impact of the negative TFP shock, and the second assesses the expected response of macroeconomic aggregates when the economy slips into a severe recession. We denote the second experiment by the Great Recession-type shock. The impact of the one-time negative aggregate shock generates a decline in consumption 0.54 percentage points larger than that of an economy with just cyclical unemployment risk. This larger decline is of a considerable magnitude. The model with just cyclical unemployment displays realistic wealth inequality and generates a decline in consumption 0.48 percentage points larger than a Representative Agent (RA) economy.[3]

In our model, the most significant fall in consumption is due to a high downside risk to labor earnings in recessions, reducing the current and expected future labor income. The Great Recession-type shock generates a sharp, deep fall in consumption and a sluggish recovery, as seen in the data. Because the Great Recession-type shock is long-lasting, it worsens household economic outlook on labor income, leading both poor and wealthy households to increase precautionary savings, reducing consumption even more. Rich households reduce consumption because of the fear of losing a high-income state. Furthermore, the high persistence of the idiosyncratic income shock makes it difficult to self-insure against it. Consequently, aggregate consumption will remain low as long as

[3]Recall that a Representative Agent economy features complete markets. In other words, a whole set of financial instruments is available to hedge all idiosyncratic risks.

households attempt to increase their precautionary savings for the following periods.

Related Literature Understanding the role of incomplete markets and household heterogeneity in the business cycle has been an active area of research since the work of Imrohoroglu (1989) and Krusell and Smith (1998).[4] This work adds to the growing literature on the relationship between wealth inequality and real macroeconomic shocks. Krueger, Mitman, and Perri (2016a), our most closely related paper, studies a model of incomplete markets with idiosyncratic income risk and preference heterogeneity to quantify how household heterogeneity, particularly wealth inequality amplifies and propagates negative real macroeconomic shocks. Their main finding is that high wealth inequality levels significantly exacerbate the aggregate consumption drop in response to a negative aggregate shock relative to the standard Representative agent economy. Wealth-poor and borrowing-constrained households with a higher marginal propensity to consume strongly cut their consumption expenditures to increase precautionary savings when the recession hits.

However, Amromin, De Nardi, and Schulze (2018) argues that KMP may underestimate the drop in consumption and the subsequent weak recovery observed in the US data because they abstract from the relevant changes during the Great Recession. In particular, and more crucial for our purposes, the KMP model assumes that unemployment risk is the only cyclical idiosyncratic risk. However, extensive and growing literature documents that in recessions, conditional on employment, large and persistent earnings declines become more likely, while increases become less likely. For instance, using labor earning data from the US Social Security Administration, Guvenen, Ozkan, and Song (2014) documented that the cyclical patterns of earnings shocks emerge from fluctuations in the skewness of the idiosyncratic earnings shock distribution rather than changes in its variance, contrary to countercyclical variance postulated by Storesletten, Telmer, and Yaron (2004b).

New evidence using large panel data sets reveals that in recessions, the right tail of the earning shock distribution collapses while the left tail enlarges, yet the median

[4]For a review on the incomplete markets research program see Krusell and Smith (2006) and Guvenen (2011)

does not change much relative to either tail, thus generating procyclical skewness. Similarly, Busch, Domeij, Guvenen, and Madera (2022) found that skewness is procyclical for continuously employed workers using administrative data from the United States, Germany, and Sweden. Changes in hours and wages are essential to generate the procyclical skewness in earnings growth. The strong procyclical skewness in earnings growth has also been documented for the UK (Angelopoulus, Lazarakis, and Malley, 2019) and Denmark (Harmerberg and Sievertsen, 2021). Moreover, Nakajima and Smirnyagin (2019) and Busch and Ludwig (2021) have documented the same pattern by modifying the parametric approach of Storesletten, Telmer, and Yaron (2004b) and using the US Panel Study of Income Dynamics (PSID).

Regarding our knowledge, McKay (2017) is the only work that studies how cyclical long-lasting income risk deepens the consumption response during recessions. He found that cyclical income shocks amplify the consumption response substantially. However, his model has a major economically counterintuitive feature: the fraction of recently unemployed workers that draw positive income shocks is larger in recessions than in expansions. In our work, we aim to correct this undesirable feature by considering that the cyclical skewness of the distribution of earning shocks causes its mean to be larger in expansions than in recessions, as suggested by Guvenen, Ozkan, and Song (2014).

This paper is organized as follows. Section 2 develops a real business cycle model with heterogeneous households, incomplete markets, and cyclical earning risk. Section 3 describes the calibration. We report the model fit to various wealth-related moments in Section 4. Then, we analyze the response of macroeconomic variables to negative aggregate technology shocks. Finally, Section 5 concludes and provides directions for future research. The Appendix contains a detailed description concerning the estimation of the stochastic process for labor earnings, its discretization, and the computational algorithms employed.

2 The 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 element is that idiosyncratic labor productivity shocks vary along the business cycle. To our knowledge, no well-established theoretical foundation explains the cyclical skewness of long-term earnings growth distribution. Therefore, we follow McKay (2017) and McKay and Reis (2021), assuming this reduced-form approach.

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) following a first-order Markov chain with transition matrix $\pi(z'|z)$. The TFP shock takes values in $\mathcal{Z} = \{z_l, z_h\}$, with $0 < z_l < 1 < z_h$. We interpret z_l as a severe recession and z_h as normal times. Let $\Pi(z)$ be the TFP shock invariant distribution. 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 newborns with zero assets, leaving the population size unchanged. The newborn households have the same individual preferences and skills as the recently deceased.

Households derive utility from consuming the final good according to a CRRA utility function with relative risk aversion parameter σ and 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, Tokuoka, and White (2017), households draw their intertemporal discount factor at the beginning of their life from a uniform distribution with support $\left[\overline{\beta} - \kappa, \overline{\beta} + \kappa\right]$.[5]

In each period, households have an endowment of one unit of time and a stochastic log-labor efficiency as the sum of a transitory shock $\nu \in \mathcal{V}$ and a persistent component $\gamma \in \mathcal{Y}$. Households supply their unit of time with labor efficiency $\exp(\nu + \gamma)$ inelastically to the labor market. Additionally, they could be either unemployed or employed. Let $s \in \mathcal{S} = \{u, e\}$ denote the current employment status of a household, with u an e denoting unemployment and employment, respectively. Employed households receive a pre-tax labor income equal to $w \exp(\nu + \gamma)$. In contrast, the unemployed receive an amount of $b = \rho w \exp(\nu + \gamma)$ from an unemployment insurance system, where $\rho \in (0, 1)$

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

represents the magnitude of the unemployment insurance generosity. Following Krueger, Mitman, and Perri (2016a), we assume that taxes are levied on both labor earnings and unemployment benefits at rate $\tau(z, \rho) \in (0, 1)$. Note that taxes may depend on the aggregate state of the economy if the unemployment rate is a function of TFP.

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 $a \in \mathcal{A} = [0, \infty)$ the household's capital or asset holdings. Capital depreciates at rate $\delta \in (0, 1)$ each period. Since households cannot borrow, markets are incomplete. Therefore, there are no financial instruments by which households can fully insure themselves against idiosyncratic risk. Consequently, households will hedge from idiosyncratic risk holding physical capital.

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

2.2.2 Idiosyncratic cyclical earnings risk

Labor earnings risk arises from two sources:

1. *Idiosyncratic unemployment risk*: in the spirit of Krusell and Smith (1998), unemployment follows a first-order Markov chain with transition matrices $\pi(s'|s, z', z)$. The matrices' dependence on the aggregate productivity transition allows the model to capture the business cycle effects on the persistence and incidence of unemployment. Because there is a continuum of agents, the Law of Large Numbers applies. Therefore, only the aggregate shock z determines the unemployment rate, which we denote by $\Pi_z(u)$.

[6]The assumption of exogenous borrowing constraints represents households' underlying frictions 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.

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

$$
\begin{aligned}
y_t &= x_t + \nu_t, \\
\text{with } \; x_t &= \varphi x_{t-1} + \eta_t
\end{aligned}
$$

The transitory component ν_t follows a normal distribution with zero mean and standard deviation σ_ν. In line with Guvenen, Ozkan, and Song (2014), the persistent component follows an AR(1) process with persistence parameter $\varphi \in (0, 1)$. The innovations of the persistent component are drawn from a mixture of normal distributions whose parameters vary with the business cycle:[9]

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

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

We use a Gauss-Hermite quadrature with n_ν nodes for the transitory innovation. Let $\sqsubseteq = \{\nu_1, \ldots, \nu_{n_\nu}\}$ be the nodes of the transitory component and let $\Pi(\nu) = \pi(\nu)$ its invariant distribution. Then, we discretize in n_γ nodes the persistent component of the log labor efficiency process. Let $\mathcal{Y} = \{\gamma_1, \ldots, \gamma_{n_\gamma}\}$ be the set of nodes employed. We assume that the discrete process follows a first-order Markov chain with transition matrices $\pi(\gamma'|\gamma, z', z)$, which depends on the aggregate state of the economy. Again, only the aggregate shock z determines the fraction of households with log-labor efficiency equal to γ, which we denote by $\Pi_z(\gamma)$. Therefore, the labor productivity of a household will be $\exp(\nu + \gamma)$.

[8] This specification finds empirical support for example in Meghir and Pistaferri (2004), Storesletten, Telmer, and Yaron (2004b), Guvenen (2009) and 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), and Guvenen, McKay, and Ryan (2022).

Without loss of generality, we make the following normalizations:

$$\sum_{\nu \in \mathcal{V}} \Pi(\nu) \exp(\nu) = 1$$

$$\sum_{z \in \mathcal{Z}} \Pi(z) \left(\sum_{\gamma \in \mathcal{Y}} \Pi_z(\gamma) \exp(\gamma) \right) = 1$$

2.2.3 Two comments on modeling idiosyncratic labor efficiency risk

In the following section, we will discuss the estimation of the labor efficiency process, some concerns, and its discretization. Nonetheless, there are some comments to be made before. First, we do not impose restrictions on the mean of the innovations of the persistent component, as the literature does. Therefore, our process's mean and median are larger in expansions than in recessions. This approach may seem flawed at first sight, but Guvenen, Ozkan, and Song (2014) argues that the cyclical nature of labor earnings 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 earnings 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 earnings shocks, with little changes in its variance (Guvenen, Ozkan, and Song, 2014). Second, due to the procyclical skewness of the idiosyncratic earnings shock distribution, if we impose some restriction on its mean, we will be assuming that households face more positive small shocks in recessions than in expansions, which is economically counterintuitive.

To illustrate the last point, consider the idiosyncratic efficiency process of Meeuwis (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 (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 (1) shows.

Figure 1: Density of annual earnings change in Meeuwis (2021).

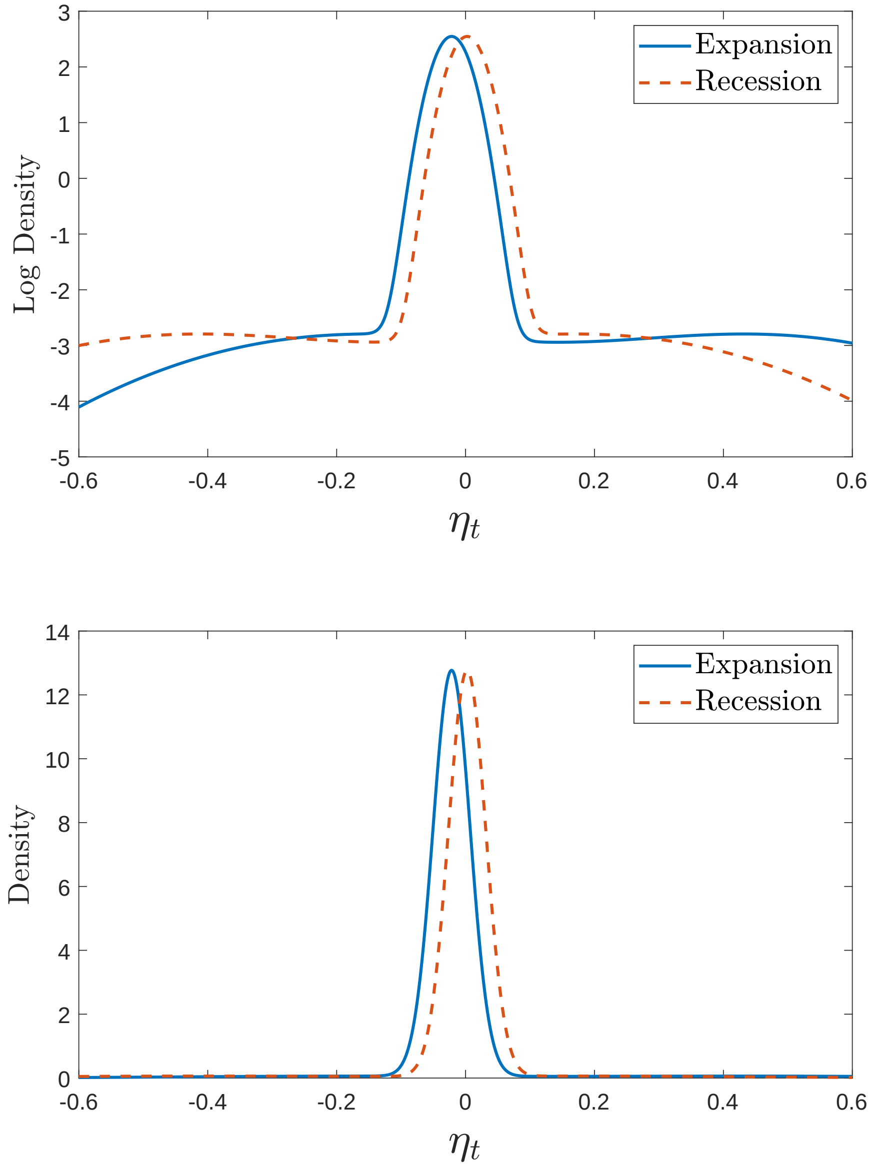

Source: own simulation using Meeuwis (2021) process and estimated parameters.

2.3 Government and social security

The Government follows a balanced budget unemployment insurance system:

$$\underbrace{\tau\left[\sum_{\nu}\Pi(\nu)\sum_{\gamma}\Pi_z(\gamma)\Big((1-\Pi_z(u))w(\Phi,z)\exp(\nu+\gamma)+\Pi_z(u)b(\nu,\gamma;\Phi,z)\Big)\right]}_{\text{Tax Revenue}}=\underbrace{\Pi_z(u)\sum_{\nu}\Pi(\nu)\sum_{\gamma}\Pi_z(\gamma)b(\nu,\gamma;\Phi,z)}_{\text{Government Spending}}$$

So, the tax rate that balances the budget is given by:

$$\tau(z,\rho)=\left(\frac{\Pi_z(u)\rho}{1-\Pi_z(u)+\Pi_z(u)\rho}\right)=\left(\frac{1}{1+\dfrac{1-\Pi_z(u)}{\Pi_z(u)\rho}}\right)$$

It is worth noting that the tax rate depends on the business cycle because there is a one-to-one mapping between the aggregate shock and the unemployment rate.

2.4 Household decision problem

Let $v(a,s,\gamma,\beta;\Phi,z)$ be the value function of a household with individual states (a,s,γ,β) when the aggregate state of the economy is (Φ,z). Therefore, the household's recursive problem is given by:

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

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

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

where $\mathbb{1}(s=u)$ is the indicator function, taking the value of one if the household is unemployed and zero otherwise. Finally, H represents the law of motion of the cross-section distribution of individual states.

2.5 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 H 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):

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

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.

3 Calibration

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

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	Krueger *et al.*, 2016a
α	0.36	Capital share	Krueger *et al.*, 2016a
ρ	50%	Unemployment replacement rate	Gruber, 1994, and Krueger *et al.*, 2016a
Business cycle parameters			
(z_l, z_h)	(0.9717, 1.0056)	Aggregate productivity shock values	Standard deviation log output: 3.1% (McKay, 2017)
$(\Pi_{z_l}(u), \Pi_{z_h}(u))$	(8.39%, 5.33%)	Unemployment rate	Krueger *et al.*, 2016a
$\pi(s'\|s, z', z)$	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			
$\overline{\beta}$	0.976	Mean discount factor	Capital to output ratio: 10.26 (Carroll *et al.*, 2017)
κ	0.007	Discount factor dispersion	Wealth Gini coefficient: 0.78
n_β	5	Number of nodes for the discretization of its distribution	Krueger *et al.*, 2016a
Persistent idiosyncratic labor efficiency shock parameters			
γ	See appendix	Idiosyncratic efficiency shock	Discretization
$\pi(\gamma'\|\gamma, z', z)$	See appendix	Transition matrix of labor efficiency process	Discretization
Transitory idiosyncratic labor efficiency shock parameter			
n_ν	3	Number of nodes for the Gaussian-Hermite quadrature	Standard value

3.1 Parameters taken from the 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. In line with Gruber (1994) and Krueger, Mitman, and Perri (2016a), wet set the unemployment replacement rate to $\rho = 50\%$. 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 $\Pi(z) = [0.164, \;\; 0.836]$.

To calibrate the values of the aggregate technology shock, we target a log-output standard deviation of 3.1% (McKay, 2017). Together with the following normalization:

$$z_l \Pi(z_l) + z_h \Pi(z_h) = 1$$

we obtained $z_l = 0.9717$ and $z_h = 1.0056$.

The idiosyncratic unemployment risk is determined by four employment-unemployment Markov transition matrices that depend on the economy's 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):

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

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

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

The parameters that characterize the distribution of the discount factor $(\overline{\beta}, \kappa)$ are calibrated to a Wealth Gini coefficient of 0.78 and a quarterly capital-to-output ratio K/Y of 10.26 (Carroll, Slacalek, Tokuoka, and White, 2017). These targeted values require that $\overline{\beta} = 0.976$ and $\kappa = 0.007$. Thus, the discount factor is uniformly distributed between $[0.9689,\ 0.9830]$. Then, the distribution is discretized in $n_{\beta} = 5$ equidistant nodes as in Krueger, Mitman, and Perri (2016a).

Incomplete market models at quarterly frequency commonly use a discount factor of approximately 0.99, and often introducing little heterogeneity in the discount factor is sufficient to match the key wealth distribution moments.[10] For instance, Krueger, Mitman, and Perri (2016a) uses discount factor values varying from 0.981 to 0.992. In our model, we also find that introducing small heterogeneity in the discount factor does the trick for matching the wealth share held by the richest quintile of the wealth distribution. However, we required discount factor values to be smaller than those commonly found in the quantitative literature to match a quarterly capital-to-output ratio of 10.26. The reason is that households in our model have strong incentives to hold an excess of precautionary wealth due to setting a larger risk aversion while facing higher risk to their current and future labor income prospects.[11] Nevertheless, the values we used are within reasonable bounds (Carroll, Slacalek, Tokuoka, and White, 2017).

3.3 The labor efficiency process

Given that our model for the labor efficiency process in section (2.2.2) is at a quarterly frequency, we search for the parameters that minimize the distance from key selected moments generated by the annual process in Guvenen, Ozkan, and Song (2014) and our quarterly process, once observations are aggregated to an annual basis. On behalf of

[10] Some examples of key wealth distribution moments that incomplete market models aim to match are the Wealth Gini coefficient, the share of the wealth that holds the richest quintile, or the share of the wealth that holds the two poorest quintiles.

[11] See Carroll, Slacalek, Tokuoka, and White (2017) for an analysis of the determinants of the mean and dispersion of the discount factor distribution.

space, we will briefly discuss the estimation method. We refer the reader to see Appendix A.3 for further information about the estimated parameters, the fitness of the estimation, and the optimization procedures employed.

In the estimation, we target the mean, median, the difference between the 50th and 10th percentile ($L5010$), the difference between the 90th and 50th percentile ($L9050$), and Kelly's measure of skewness of 1, 3, and 5-year income changes. Also, we target the autocorrelation of the annual process. These moments capture how the distribution of labor income shocks changes over the business cycles. Kelly's skewness is of particular interest, and to see why let us look at its definition:

$$\mathcal{K} = \frac{L9050 - L5010}{L9010}$$

It can be seen in the definition that the interpretation of Kelly's skewness is straightforward. It measures how much of the overall dispersion $L9010$ is due to the left tail dispersion $L5010$ and the right tail dispersion $L9050$. Therefore, it allows us to capture how the tails shrink and widen over the business cycle. For instance, in recessions, the lower tail enlarges, increasing the difference $L5010$, and the upper tail shrinks, decreasing the difference $L9050$, yielding negative skewness.

3.4 Addressing some concerns regarding the labor efficiency process

Because most studies use annual data, one possible concern of the cyclical sources of labor income earning changes is that nonemployment periods could generate procyclical skewness. However, Guvenen, McKay, and Ryan (2022) estimated an income process featuring procyclical skewness, and they found that the percentile-based measure of skewness is not affected by nonemployment periods. Similarly, McKay (2017) arrives at the same conclusion. Nevertheless, these results do not hold for the third central moment, which is largely affected by nonemployment periods (McKay, 2017). Additionally, it is worth noting that unemployment periods can explain almost entirely the leptokurtic nature of the income growth distribution (Guvenen, McKay, and Ryan, 2022).

4 Results

The results are based on a comparison of two versions of the model. The first version assumes that the labor efficiency process follows an AR(1) with innovations drawn from a mixture of normal distributions with constant parameter values. This version approximates the original KMP model without the life-cycle component and different parameterization. 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 whose parameters vary along the business cycle, so the distribution of labor earning growth exhibits a procyclical skewness. Thus, in this model version, conditional on employment, long-lasting declines in labor earnings are more likely during recessions, while large, long-lasting upward movements are less likely. We denote this version as the Cyclical Model.

The comparison of the models is in terms of their ability to match the empirical US cross-sectional wealth distribution, the aggregate consumption, investment, and output response to negative technology shocks, and the consumption response across the wealth distribution. Regarding the response to negative technology shocks, we evaluate two experiments. The first experiment assesses the economy's response to a one-time negative technology shock. The second experiment computes the expected response to a stochastic duration-type negative technology shock lasting 22 quarters on average. To make the comparisons fair between models, they display the same initial output drop, the same average capital-to-output ratio, and the Gini coefficient for wealth distribution when the economy falls into recession.

4.1 Model Fit

The key elements that allow the KMP model to replicate the empirical distribution of wealth are the heterogeneity in the discount factor and the idiosyncratic efficiency risk. Heterogeneity in the degree of household patience for future consumption streams allows a non-negligible share of extremely patient households to continue saving even at high levels of wealth. At the same time, it produces very impatient households with little incentive to accumulate wealth, amplifying wealth inequality. Moreover, 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, making it more difficult for them to accumulate wealth. Conversely, households in a high-productivity state will accumulate wealth for fear of a negative labor efficiency shock. Given those mentioned above, we explore to what extent the inclusion of persistent and cyclical income shocks, conditional on employment, affects the model's ability to replicate the observed US cross-sectional wealth distribution.

4.1.1 Matching the Wealth Distribution

Table (2) reports the share of wealth accumulated by wealth quintiles calculated from the Survey of Consumer Finances (SCF) of 2007 and our two model versions.[12] The table shows several interesting facts. First, the table presents negative entries that correspond to debt. Since borrowing is forbidden in our models, we aim to match as closely as possible the fact that, added together, the first two quintiles hold almost no wealth. Our models fit the empirical distribution of wealth in the US quite well at the bottom 40% but worse at the top. In our model versions, the first two quintiles account for 1.4% of total wealth. Albeit this fit is not perfect, it is similar to what the KMP model reports. Overall our model fits reasonably well the concentration of wealth held by all quintiles, doing slightly worse in the richest quintile. However, as is common in these types of models, it has trouble replicating the wealth share of the richest 10%, 5%, and especially the 1%. Nevertheless, we do not consider this a major drawback because the consumption policy function is almost linear at high wealth levels. Therefore, we may think that mechanically redistributing wealth beyond the top wealthiest 20% does not significantly alter aggregate consumption dynamics. Second, as the table (2) shows, the distribution of wealth is virtually identical in our two models. Therefore, the inclusion of countercyclical income risk

[12] PSID 2006 and SCF 2007 wealth distribution statistics are taken from Krueger, Mitman, and Perri (2016a).

does not significantly change the distribution of wealth generated by the Acyclical Model.

Figure (2) presents the Lorenz curve for the wealth distribution from the SCF 2007 data and the Cyclical Model.[13] The figure displays the patterns documented in the previous paragraph. Our models have difficulty fitting the wealth distribution at the top but match the wealth distribution's bottom and middle portions remarkably well.

Table 2: Wealth Distribution: Data v/s Model

	Data		Models	
% Share held by:	PSID, 06	SCF, 07	Acyclical	Cyclical
Q1	-0.9	-0.2	0.2	0.2
Q2	0.8	1.2	1.2	1.2
Q3	4.4	4.6	4.2	4.2
Q4	13.0	11.9	13.2	13.3
Q5	82.7	82.5	81.3	81.2
90−95	13.7	11.1	16.2	16.0
95−99	22.8	25.3	27.2	27.2
Top 1%	30.9	33.5	21.2	21.2
Gini	0.77	0.78	0.78	0.78

Note: This table reports wealth distribution statistics computed from the Panel Study of Income Dynamics (PSID) 2006 and the Survey of Consumer Finances (SCF) 2007 data, all taken from Krueger, Mitman, and Perri (2016a).

Figure 2: Lorenz Curve: SCF, 07 v/s Model

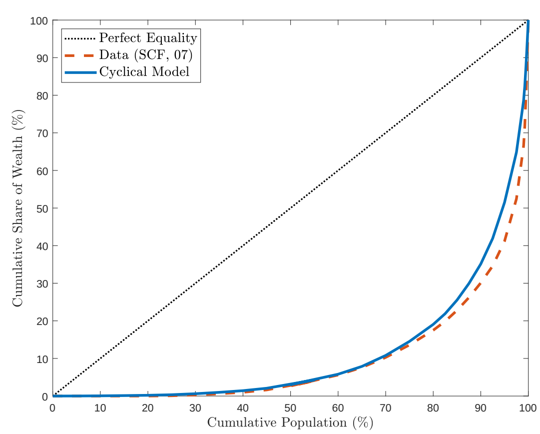

Note: This figure displays the wealth distribution Lorenz Curve for SCF 07 data and the Cyclical Model.

[13]We do not include the Lorenz curve for the Acyclical Model because it is virtually identical to its cyclical counterpart.

4.2 The Dynamics of Macroeconomic Aggregates in Severe Recessions

The main finding of Krueger, Mitman, and Perri (2016a) is that an incomplete market economy that generates realistic wealth heterogeneity amplifies the initial drop in aggregate consumption by approximately 25% when the economy slips into a recession relative to a low-wealth heterogeneity economy. The reason is that the KMP model generates a wealth distribution characterized by 40% of the population holding no wealth while accounting for a significant share of aggregate consumption. Hence, when a negative aggregate productivity shock hits the economy, it increases the probability of unemployment, so these wealth-poor households will strongly reduce consumption to increase their precautionary savings. Because unemployment spells are short, agents face an increased likelihood of a low-persistence negative shock in recessions. Nevertheless, due to incomplete markets, the lack of full insurance generates a stronger reduction in consumption than in the low-wealth heterogeneity economy.

Introducing a higher likelihood of long-lasting declines in income while reducing the probability of upward movements for employed households when the economy slips into a recession should strengthen the precautionary savings motive, amplifying the aggregate consumption drop and weakening its subsequent recovery (McKay, 2017; Amromin, De Nardi, and Schulze, 2018). In this subsection, we provide a quantitative answer to this hypothesis. Following Krueger, Mitman, and Perri (2016a), we consider two quantitative experiments. In both experiments, we take as an initial condition the distribution of wealth produced after several realizations of normal times for the aggregate productivity shock, so the wealth distribution is stabilized. Then, the economy slips into a severe recession. The recession lasts only one quarter in the first experiment, returning to normal times afterward. In the second experiment, the economy goes into recession for one quarter, and after that, it evolves stochastically, according to its aggregate shock transition matrix. Therefore, in the second experiment, the expected duration of the recession is 22 quarters.[14] We simulate 10,000 independent paths of aggregate productivity shocks.

[14]Note that this experiment is different from the one performed in Krueger, Mitman, and Perri (2016a). They study the response of macro aggregates when the economy goes into a recession that lasts 22 quarters,

Then, for each period, we average across simulations the responses of the macroeconomic variables.

4.2.1 A One-Time Negative Technology Shock

To aid the comparison in the one-time negative technology shock, we add the response of an RA economy. Figure (3) 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 RA and Acyclical Model to match the same initial output drop in recessions that generates the Cyclical Model.[15] The figure (3) 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 RA economy. Thus, the same output decline generates 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 RA 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.

and then the economy returns to normal times.

[15] Recall that all models are calibrated to match the same capital to output ratio. Also, the Acyclical model is calibrated to match the same wealth Gini coefficient of the Cyclical model.

Figure 3: Impulse Response: one-time negative technology shock

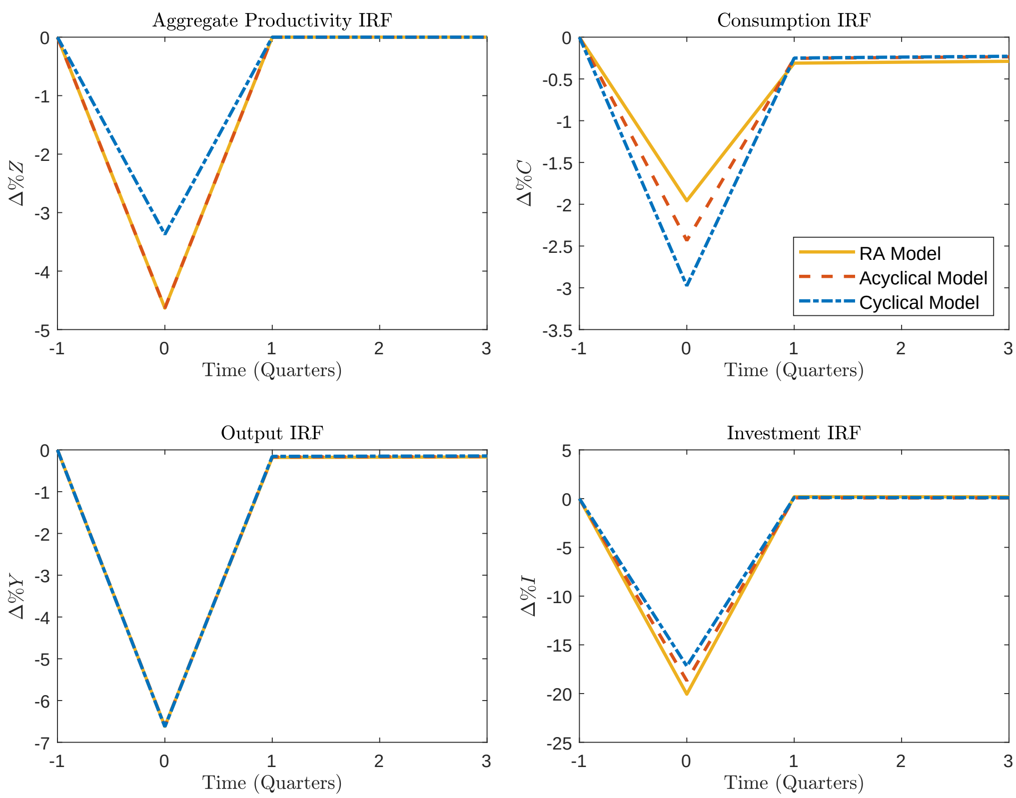

Note: The figure displays consumption, investment, and output dynamics response to a one-time negative technology shock.

4.2.2 The impulse response function of consumption across wealth quintiles

In the previous paragraph, we have documented that the difference between the consumption drop in the Acyclical and the Cyclical model was due to the increased probability of suffering a long-lived idiosyncratic negative income shock in recessions. However, the figure (3) could mask stronger consumption drops when looking at different parts of the wealth distribution. The previous thought is verified by the table (3), which shows the consumption drop by wealth quintile for both model versions and the percentage point difference between them. Table (3) shows that the percentage point difference has a U shape, which could be explained by the swings of the tails of the idiosyncratic efficiency shock distribution. Because efficiency is highly persistent, its correlation with net worth is significant, so the poorest and richest quintiles will be the most affected by the idiosyncratic efficiency risk when the economy slips into a recession.

Table 3: Consumption drop by Wealth Quintile

Wealth Quintile	Acyclical Model	Cyclical Model	Δ (pp)
Q1	-5.69%	-6.15%	0.46
Q2	-3.70%	-4.12%	0.42
Q3	-3.01%	-3.43%	0.42
Q4	-2.60%	-3.02%	0.42
Q5	-1.94%	-2.57%	0.63

Note: this table shows the consumption drop by wealth quintile and model. It also shows the percentage point difference between models by wealth quintile.

Why is the consumption drop so large for the first two wealth quintiles, while the aggregate consumption drop is not so strong? The explanation is that the first two wealth quintiles account for a small share of aggregate consumption. Table (4) shows the share of disposable income (labor plus capital income rents plus unemployment insurance) and expenditures (consumption) by wealth quintile. As shown by the table (4), in the Cyclical Model, the first two wealth quintiles account for approximately 11% of aggregate consumption, while it is 23% in the data, which is explained by the low share of disposable income of the first two wealth quintiles. Because the model fails to match the joint distribution of wealth and disposable income, it will also fail to match the joint distribution of wealth and consumption.

So why does the model fail to match the joint distribution of wealth and disposable income? The answer relies on two factors. In the first place, the variance of the idiosyncratic efficiency process is 38% larger than in KMP. Secondly, the process displays high autocorrelation. The first factor is no surprise given that the quarterly idiosyncratic efficiency process parameters are obtained employing the estimated parameters of Guvenen, Ozkan, and Song (2014), which use a non-capped database. Because there is a high correlation between wealth and the magnitude of the households' efficiency state, the share of

Table 4: Selected Variables by Net Worth: Data v/s Model

	% Share of			
	Disposable Income		Expenditures	
Net Worth	PSID, 06	Cyclical Model	PSID, 06	Cyclical Model
Q1	8.7	3.2	11.3	2.9
Q2	11.2	7.2	12.4	6.9
Q3	16.7	12.2	16.8	11.6
Q4	22.1	20.6	22.4	19.9
Q5	41.2	56.8	37.2	58.7

Note: This table reports the share of disposable income and consumption expenditure by net worth computed from the PSID 06 data and the Cyclical Model. PSID joint distribution statistics are from Krueger, Mitman, and Perri (2016a).

consumption by the poor-wealth households will be lower than the share observed in the data. A straightforward solution is incorporating progressive taxation as in Heathcote, Storesletten, and Violante (2017). We outline how to include a progressive tax system in the appendix (A.2). In this proposed solution, we assume that incorporating a progressive tax-transfer system that shifts resources from the wealthiest to the poorest households would increase the consumption share of the poor wealth, which was confirmed in exercises (not reported).

4.2.3 Expected Severe Recession-Type Shock

Figure (4) 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.

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.

Considering both IRFs experiments, it is noteworthy to mention that our model generates a consumption drop larger and a languish recovery after the initial drop than the model in McKay (2017). In the next section, we explore the role of TFP in the dynamics of consumption in response to negative technology shocks.

4.2.4 Response with and without TFP

What is the role of TFP changes in the IRFs we have analyzed? Recall that the model of McKay (2017) falls short in accounting for the deep decline in consumption observed during the Great Recession because, in his results, TFP remains constant. Notice that considering TFP changes in the last experiment, our model can match the magnitude of the consumption drop observed in the data of approximately 3.6%. To gain further insights into TFP changes' role, we repeat the two experiments in an economy similar to the Cyclical Model but keeping TFP constant in recessions and expansions. The figure (5) shows the consumption response for the Cyclical Model with and without TFP changes for both

Figure 4: Impulse Response: Severe recession technology shock.

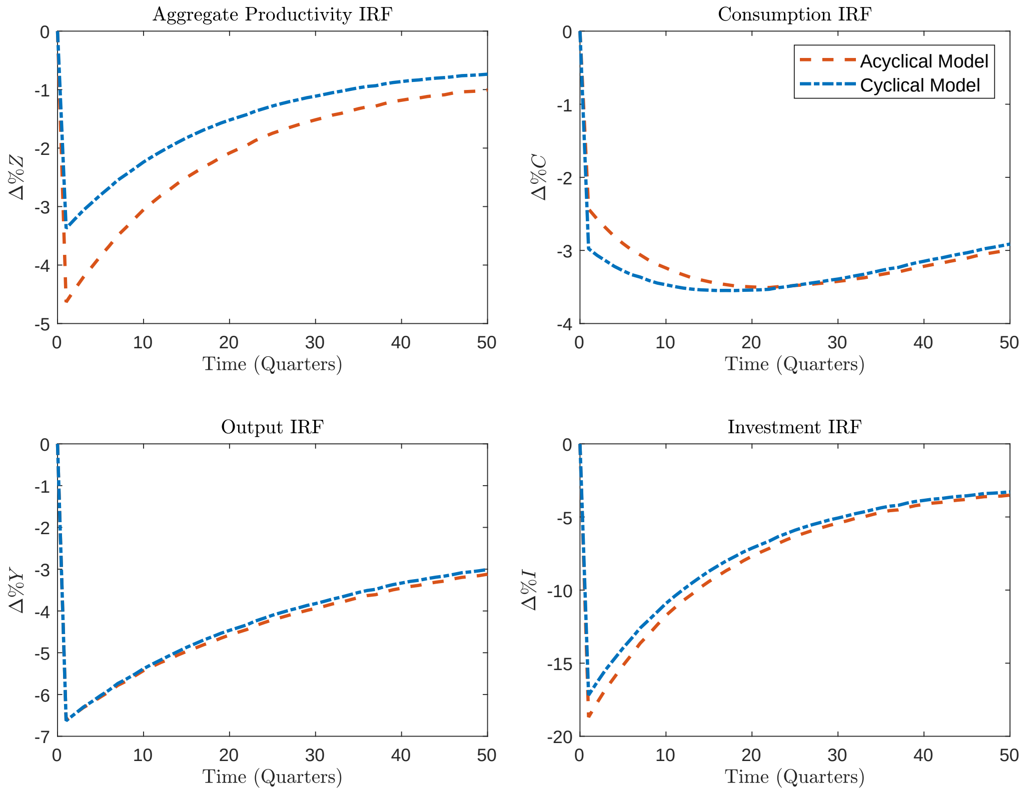

Note: The figure displays the expected consumption, investment, and output response dynamics when the economy slips into a recession.

experiments.

We want to highlight three observations from the figure (5). First, lower factor prices when the recession hits the economy is crucial for generating a sharp initial consumption drop. The intuition behind the further initial consumption fall is that when TFP is lower, the disposable income of households is decreased because the real interest rate and real wage decrease, reducing the net return of asset holdings and labor income of all households, including for those who do not draw an idiosyncratic negative shock because of the recession. Second, in response to a Severe Recession Technology Shock, the model with constant TFP displays negligible additional falls in consumption after the initial drop. In contrast, in the model where TFP decreases, consumption falls an additional 0.5 per-

Figure 5: Consumption Response: One-Time and Severe Recession Technology Shock for the Cyclical Model with and without TFP changes.

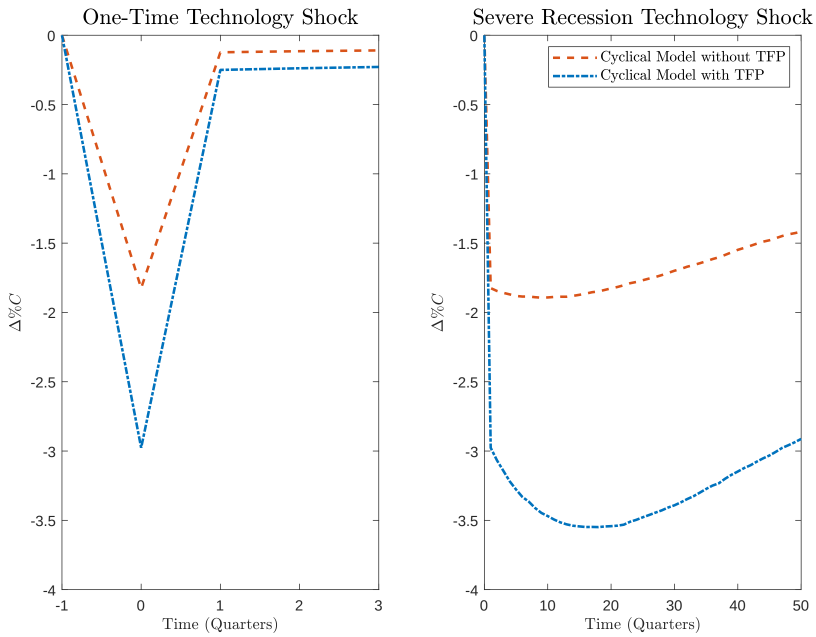

Note: The figure displays the dynamics of the consumption response when the economy slips into a recession that last one period and a recession that lasts, on average, 22 periods.

centage points after the initial drop. Third, in response to a Severe Recession Technology Shock, consumption recovers faster in the model with constant TFP, where it begins to recover at period 10, relative to the model with TFP changes, where it starts to grow after 18 periods. The slow recovery when TFP falls is because households will allocate fewer goods for investment and consumption due to the reduction in available sources. Therefore, for consumption to begin to recover, households need to ensure that they have reached an optimal level of savings (investment), which takes more time when recessions are associated with lower TFP.

Given the observations in the previous paragraph, we conclude that cyclical idiosyncratic efficiency shocks are an important element to consider in generating an initial sharp decline in consumption and a weaker recovery, even if TFP remains constant. However, cyclical idiosyncratic efficiency shocks and reductions in TFP reinforce the initial sharp drop in consumption and slow its recovery, making the consumption response in the model closer to the response shown by the data. Consequently, future research studying severe recessions cannot leave aside changes in TFP. In this regard, McKay (2017) acknowledges that his results merely illustrate the importance of cyclical long-term earnings risk rather than fully characterize the consumption dynamics during the Great Recession, that is, its sharp initial drop and its languished recovery.

5 Conclusion

This paper adds to a growing literature emphasizing the importance of countercyclical earning risks during recessions for consumption dynamics. We have argued that the inclusion of cyclical 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 0.54 percentage points 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 lead households to sharply cut consumption and increase their precautionary savings to insure themselves against the possibility of suffering a persistent fall in earnings during economic downturns.

In this work, the more consumption drops, the faster the recovery from recessions is. Moreover, social insurance has no role in this paper 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 can be taken to shed light on the importance of public policies employed during severe recessions.

First, as Krueger, Mitman, and Perri (2016a) did, aggregate demand externalities could model the negative loop between output and consumption. A drop in consumption would yield an additional output reduction, lowering aggregate wages and further exacerbating the consumption drop. Thus, social insurance programs aiming to reduce wealth inequality stabilize consumption, decreasing the business cycle fluctuations.

Second, in our model, the first two wealth quintiles account for 10% of aggregate consumption, while it is 23% in the data, which happens because the model generates a low share of earnings for the first two wealth quintiles. Incorporating progressive taxation as in Heathcote, Storesletten, and Violante (2017) should fix this issue.[16] Because low-wealth households would receive larger transfers, the consumption drop due to precautionary saving motives will be strengthened due to low-wealth households having a higher marginal propensity to consume than the wealthiest.

[16]See Appendix A.2 for the details on how to implement a progressive labor income tax system.

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).

Arnoud, A., Guvenen, F., and Kleineberg T. (2019). Benchmarking Global Optimizers. *Working Paper*.

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.

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., Tokuoka, 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.

Chaterine, S. (2021). Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle. *The Review of Financial Studies*.

Chetty, R. (2008). Moral Hazard versus Liquidity and Optimal Unemployment Insurance. *Journal of Political Economy* 116(2), 173-234.

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 Working 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.

Gruber, J. (1994). The Consumption Smoothing Benefits of Unemployment Insurance. *NBER Working Paper Series*, Working Paper 4750.

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). Macroeconomics 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.

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

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

Heathcote, J., Storesletten, K. and Violante, G. (2017). Optimal Tax Progressivity: an Analytical Framework. *Quarterly Journal of Economics* 132(4), 1693-1754.

Hubmer, J. (2018). The job ladder and its implications for earnings risk. *Review of Economic Dynamics* 29, 172-194.

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. (1989). Cost of Business Cycles with Indivisibilities and Liquidity Constraints. *Journal of Political Economy*, 97(6), 1364-1383.

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

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

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.

Lkhagvasuren, D., and Bataa E. (2022). Finite-State Markov Chains with Flexible Distributions. *Computational Economics*.

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*.

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. (2020). Skewed Business Cycles. *Quarterly Journal of Economics*, R&R.

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 Over the Life Cycle. *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.

A Appendix

A.1 Krusell and Smith's Quasi-Aggregation algorithm

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).[17,18] 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 $(a, s, \nu, \gamma, \beta)$ solves the following recursive problem:

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

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

where ψ_l, ψ_h, κ_l and κ_h are constants to be determined using the Krusell and Smith (1998)

[17] 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.

[18] See Appendix A.5 for details on the algorithm employed to simulate a continuum of agents.

method. We iterate on the Euler equation to solve the household decision problem, as in Maliar, Maliar, and Valli (2010).[19]

A.2 Extension: Progressive Labor Income Tax System

Following Heathcote, Storesletten, and Violante (2017), the net tax on labor income for an employed household with idiosyncratic efficiency (ν, γ) is

$$w(K,z)\exp(\nu+\gamma) - (1-\tau)w(K,z)\exp\Big((1-\zeta)(\nu+\gamma)\Big)$$

The parameter τ determines the tax level, and the parameter ζ controls the tax system's progressivity. On the one hand, if there is no progressivity, that is, $b = 0$, the tax system is linear, recovering the same tax schedule as the initial model. On the other hand, if $b = 1$, the magnitude of the progressivity is so large that all households end with the same after-tax labor income. A balanced government budget implies that:

$$\underbrace{\Big(1-\Pi_z(u)+\rho\Pi_z(u)\Big)\sum_{\nu}\Pi(\nu)\sum_{\gamma}\Pi_z(\gamma)\Bigg[w(z,K)\exp(\nu+\gamma)-(1-\tau)w(K,z)\exp\Big((1-\zeta)(\nu+\gamma)\Big)\Bigg]}_{\text{Tax Revenue}} =$$

$$\underbrace{\rho\Pi_z(u)\Bigg[\sum_{\nu}\Pi(\nu)\sum_{\gamma}\Pi_z(\gamma)w(K,z)\exp(\nu+\gamma)\Bigg]}_{\text{Government Spending}}$$

Using that $\sum_{\nu}\Pi(\nu)\exp(\nu) = 1$ and after some algebra we obtain:

$$1-\tau(z,\rho,\zeta) = \frac{\Bigg(1-\Pi_z(u)\Bigg)\Bigg[\sum_{\gamma}\Pi_z(\gamma)\exp(\gamma)\Bigg]}{\Bigg(1-\Pi_z(u)+\rho\Pi_z(u)\Bigg)\Bigg[\sum_{\nu}\Pi(\nu)\sum_{\gamma}\Pi_z(\gamma)\exp\Big((1-\zeta)(\nu+\gamma)\Big)\Bigg]}$$

Therefore, after-tax income for a household with efficiency (ν, γ) is:

$$\Big(1-\tau(z,\rho,\zeta)\Big)w(K,z)\exp\Big((1-\zeta)(\nu+\gamma)\Big)\Big[1-(1-\rho)\mathbb{1}(s=u)\Big]$$

Finally, the model can be solved by replacing the new tax level and setting $\zeta = 0.181$ (Heathcote, Storesletten, and Violante, 2017).

[19]See Appendix A.6 for details of the Euler equation method.

A.3 Estimation of the Labor 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 a process of log-labor productivity shocks. The estimated process features parameters that vary over the business cycle. Given that our quantitative model is at a quarterly frequency, we translate the annual process to a quarterly process by minimizing the distance from selected moments once quarterly observations are aggregated. Those moments try to capture how the labor income distribution shocks change over the business cycles, especially how, in recessions, the right tail collapses while the left tail enlarges.

The procedure targets the mean, median, and difference between the 50th and 10th percentile ($L5010$), between the 90th and 50th percentile ($L9050$), and Kelly's measure of skewness of 1, 3, and 5-year income changes. We also target the autocorrelation of the annual process.

The quarterly process proposed has the same structure as its annual counterpart:

$$\begin{aligned} y_t &= x_t + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon) \\ \text{and} \;\; x_t &= \varphi x_{t-1} + \eta_t \end{aligned}$$

where η_t follows a mixture of normal distributions:

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

Note that the mixture's means and probabilities change along with the economy's aggregate state.

The optimization procedure we used to find the parameter values for the quarterly process is described below:

1. We simulate a long time series of aggregate shocks using the matrix outlined in section 3. Then, we obtain annual observations by simulating a large panel of earning histories using the annual process parameters and compute the mean, median, 10th, 50th, and 90th percentile of 1, 3, and 5-year changes and Kelly's measure of skewness.

2. Then, we use a global optimization procedure and a local optimization algorithm to find the quarterly process's parameter values that minimize the percentage difference between the moments generated by the annual process and the same moments obtained through aggregating the simulated quarterly observations. To search the parameter space extensively, we use Matlab's built-in Particle Swarm Optimization (PSO) algorithm.[20] Then, for each vector of parameters obtained from the PSO algorithm, we use Nelder-Mead's downhill simplex algorithm.[21,22]

The estimated parameters are the following:

Table A.1: Estimated parameters for the quarterly labor efficiency process:

φ	$p_{1,R}$	$p_{1,E}$	$\mu_{1,R}$	$\mu_{2,R}$	$\mu_{1,E}$	$\mu_{2,E}$	σ_1	σ_2	σ_ϵ
0.994	0.247	0.291	-0.171	0.049	0.115	-0.044	0.178	0.007	0.088

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

[21] See Arnoud, Guvenen, and Kleineberg (2019), and Appendix D.3 of Guvenen, Karahan, Ozkan, and Song (2021) for a detailed explanation of global optimization algorithms used in non-smooth optimization problems.

[22] We also tried Matlab's built-in Genetic algorithm optimization, but our results barely changed.

A.4 Discretization of an AR(1) with Non-Gaussian Innovations

To discretize the AR(1) with Non-Gaussian innovations, we proceed as follows:

1. Use the method outlined in Lkhagvasuren and Bataa (2022) to discretize the acyclical version of the process in section (2.2.2). Obtain the nodes $\{\gamma_1, \ldots, \gamma_N\}$, the transition matrix $\pi(\gamma'|\gamma)$ and the respective invariant distribution $\Pi(\gamma)$.

2. Simulate a long series for the aggregate shock and a long and large panel of households using the estimated parameters for the quarterly process. Denote this panel of observations by $\{x_{i,t}\}_{i=1,t=1}^{I,T}$, where I is the number of agents and T is the length of the simulation. Then, discretize the continuous simulated observations by choosing the nearest node obtained in step 1. Denote this panel of discrete observations by $\{\gamma_{i,t}\}_{i=1,t=1}^{I,T}$.

3. Compute the long-term distribution of the discrete process for recession $\Pi_{z_l}(\gamma)$ and expansions $\Pi_{z_h}(\gamma)$, as follows:

$$\Pi_z(\gamma_j) = \left(\frac{1}{IT}\right)\left(\sum_{t=1}^{T}\left[\sum_{n=1}^{I}\mathbb{1}\left(\gamma_{i,t} = \gamma_j, z_t = z\right)\right]\right)$$

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

4. Make educated guesses of the transition matrix for $\pi(\gamma'|\gamma, z, z')$. A good guess to start with is the transition matrix obtained for the acyclical process. Denote the initial guesses by $\pi^0(\gamma'|\gamma, z, z')$.

5. On the jth iteration, $\forall(z, z') \in \mathcal{Z} \times \mathcal{Z}$, obtain $\lambda_j, \quad j = 1, \ldots, n$ from:

$$\begin{pmatrix} \Pi_{z'}(\gamma_1) \\ \vdots \\ \Pi_{z'}(\gamma_N) \end{pmatrix} = \begin{pmatrix} \lambda_1 \pi^j(\gamma_1|\gamma_1, z, z') & \ldots & \lambda_1 \pi^j(\gamma_N|\gamma_1, z, z') \\ \vdots & \ddots & \vdots \\ \lambda_N \pi^j(\gamma_1|\gamma_N, z, z') & \ldots & \lambda_N \pi^j(\gamma_N|\gamma_N, z, z') \end{pmatrix} \begin{pmatrix} \Pi_z(\gamma_1) \\ \vdots \\ \Pi_z(\gamma_N) \end{pmatrix}$$

Then, normalize the probabilities to sum 1:

$$\sum_{k=1}^{N} \pi^{j+1}(\gamma_i|\gamma_k, z, z') = 1, \;\; \forall\, i = 1, \ldots, N, \;\; \forall (z, z) \in \mathcal{Z} \times \mathcal{Z}$$

6. Use π^{j+1} as the next guess. Iterate until convergence:

$$\left|\left|\pi^{j+1}(\gamma'|\gamma, z, z') - \pi^{j}(\gamma'|\gamma, z, z')\right|\right|_{\max} < 10^{-7}, \;\; \forall (z, z) \in \mathcal{Z} \times \mathcal{Z}$$

A.4.1 Illustrating the results of the discretization

To illustrate the increased risk when the economy slips into an economic downturn we present the transition matrices resulting from the discretization of the efficiency shock when the aggregate state of the economy remains in normal times and when it is transitioning from normal times to recession.

$$\pi\left(\gamma'|\gamma, z_h', z_h\right) = \begin{pmatrix}
0.9790 & 0.0207 & 0.0002 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0.0032 & 0.9739 & 0.0226 & 0.0002 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0.0052 & 0.9741 & 0.0204 & 0.0002 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0.0077 & 0.9748 & 0.0174 & 0.0001 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0.0105 & 0.9750 & 0.0143 & 0.0001 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0.0001 & 0.0138 & 0.9746 & 0.0115 & 0.0001 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0.0001 & 0.0173 & 0.9736 & 0.0090 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0.0002 & 0.0208 & 0.9723 & 0.0068 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0.0003 & 0.0237 & 0.9712 & 0.0049 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0003 & 0.0250 & 0.9715 & 0.0032 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0003 & 0.0214 & 0.9783
\end{pmatrix}$$

$$\pi\left(\gamma'|\gamma, z_l', z_h\right) = \begin{pmatrix}
0.9973 & {\color{blue}0.0027} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
{\color{red}0.0252} & 0.9702 & {\color{blue}0.0046} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
{\color{red}0.0006} & {\color{red}0.0253} & 0.9682 & {\color{blue}0.0060} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & {\color{red}0.0004} & {\color{red}0.0259} & 0.9670 & {\color{blue}0.0067} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & {\color{red}0.0003} & {\color{red}0.0268} & 0.9663 & {\color{blue}0.0065} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & {\color{red}0.0004} & {\color{red}0.0296} & 0.9645 & {\color{blue}0.0054} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & {\color{red}0.0005} & {\color{red}0.0356} & 0.9599 & {\color{blue}0.0040} & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & {\color{red}0.0009} & {\color{red}0.0448} & 0.9518 & {\color{blue}0.0025} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & {\color{red}0.0014} & {\color{red}0.0598} & 0.9374 & {\color{blue}0.0014} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & {\color{red}0.0026} & {\color{red}0.0798} & 0.9171 & {\color{blue}0.0006} \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & {\color{red}0.0044} & {\color{red}0.0992} & 0.8964
\end{pmatrix}$$

The red probabilities indicate the probability of falling to a state with lower idiosyncratic efficiency, and the blue probabilities indicate the probability of rising to a state with higher idiosyncratic efficiency. It is noteworthy to notice (i) how dramatically the probability of falling to a state with lower idiosyncratic efficiency increases and of an upward movement decreases when the economy slips from normal times to recession. This rise in the likelihood of negative income shocks and decrease in the likelihood of positive income shocks could be an important mechanism to take into consideration when accounting for the drop in aggregate consumption and (ii) how the changes in the magnitude of the probabilities at the end of the distribution of the discrete process.

A.5 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
   $$\omega \in \Omega = \{u, e\} \times \{\nu_1, \ldots, \nu_{n_\nu}\} \times \{\gamma_1, \ldots, \gamma_{n_\gamma}\} \times \{\beta_1, \ldots, \beta_{n_\beta}\}$$
   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:
   $$a'(x_{\omega,i,t}, s, \nu, \gamma, \beta; K_t, z_t) = \kappa_i$$

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.6 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:

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

Form the budget constraint:

$$c(a', s, \nu, \gamma) = \left[\frac{1-\delta+r}{\theta}\right] a + (1-\tau) w \exp(\nu+\gamma)\Big[1 - (1-\rho)\mathbb{1}(s=u)\Big] - a'$$

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

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

$$
\begin{aligned}
\Leftrightarrow \quad \tilde{a}'(s) = {} & \left[\frac{1-\delta+r}{\theta}\right] a + (1-\tau) w \exp(\nu+\gamma)\Big[1 - (1-\rho)\mathbb{1}(s=u)\Big] \\
& - \left\{ h + \beta \mathbb{E}\left[\frac{1-\delta+r'}{\left(\left[\frac{1-\delta+r'}{\theta}\right] a' + (1-\tau') w' \exp(\nu'+\gamma')\Big[1 - (1-\rho)\mathbb{1}(s'=u)\Big] - a''\right)^{\sigma}}\right]\right\}^{-\frac{1}{\sigma}}
\end{aligned} \tag{2}
$$

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

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 $a'(a, s, \nu, \gamma; K, z)$ for the values on the grid.

3. For each point in the grid $(a, s, \nu, \gamma; K, z)$, plug the policy function $a'(a, s, \nu, \gamma; K, z)$ on the right side of equation (2), set the Lagrangian multiplier to equal zero, and compute the new policy function for capital, $\tilde{a}'(a, s, \nu, \gamma; K, z)$. For any point in the grid such that $\tilde{a}'(a, s, \nu, \gamma; K, z)$ is not in the range $[0, a_{max}]$, set $\tilde{a}'(a, s, \nu, \gamma; K, z)$ equal to the value of the corresponding limit.

4. Update the policy function using the following formula:

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

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