Dynamics of the Wealth Distribution in the Presence of Higher-Order Earnings Risk

Leon Huetsch

University of Pennsylvania

October 11, 2023

Abstract

This paper introduces higher-order earnings risk consistent with recent empirical findings into a benchmark heterogeneous-agent macro model to examine its implication for the distribution of wealth. I find that higher-order earnings dynamics induce higher earnings inequality driven primarily by persistent earnings losses at the bottom. Poor households respond by strongly cutting consumption leading to more consumption and less wealth inequality which reinforces the known issue of generating the empirically observed wealth dispersion in this class of models. In addition to lower overall consumption, the higher-order earnings moments, particularly excess kurtosis, are passed through to consumption dynamics of the poor. Both effects combined mean that those households are willing to pay up to 1.7% of permanent consumption to avoid higher-order earnings risk. Moreover, the latter effect induces consumption dynamics of the poor to be predominantly driven by idiosyncratic earnings changes which significantly reduces the correlation between their consumption and aggregate output. Since wealthier households are not affected strongly the implications for the aggregate dynamics of the economy are negligible. Methodologically, I develop a new General Polynomial Chaos Expansion approach to solve for the aggregate dynamics of this class of models, and contrast its efficiency with previous methods.

1 Introduction

The aim of this paper is to explore the long and short-term dynamics of the wealth distribution in the presence of higher-order earnings risk. A growing body of recent empirical studies using administrative and survey panel data on individual earnings finds that earnings dynamics are richer than usually assumed and modeled in quantitative macroeconomic models. Specifically, these studies document that the distribution of shocks to individual earnings exhibits sizable left-skewness and substantial excess kurtosis. This is in contrast to most standard approaches of capturing individual earnings risk. The most commonly used example is the canonical linear transitory plus persistence process with a Gaussian distribution for both idiosyncratic shocks which can neither capture any higher-order moments in the distribution of shocks to individual earnings nor any dependence of persistence or moments on the earnings history.

Earnings dynamics, coupled with the net value of asset holdings, play a central role in determining consumption responses to earnings risk and shocks over the life- and business cycle. Individuals save precautionarily when facing a high degree of earnings risk in order to at least partially insure against potential future changes in earnings and they respond to the realization of an unexpected change in earnings by altering consumption behavior. For both, the precautionary savings motive as well as the consumption response to shocks, the size and persistence of earnings changes matters in determining how much consumption and saving behavior needs to adjust to ensure a certain standard of living today and in the future. In line with that, several recent studies, among others Busch and Ludwig (2020) and De Nardi et al. (2018), show that the existence of higher-order moments in the distribution of earnings changes has important implications for consumption and saving choices. Based on these findings, I explore how the presence of higher-order earnings risk affects different parts of the cross-sectional household distribution and what the resulting equilibrium implications for long and short-run wealth inequality are. I then evaluate the implications of these findings for the ability of standard heterogeneous-agent models to match the observed degree of wealth inequality.

The main contribution of this paper is two-fold. First, I evaluate the role of higher-order earnings risk for consumption, earnings and wealth in the cross-section and over the business cycle. In order to do so I use an incomplete markets real business cycle model in which households face aggregate and idiosyncratic income risk, and accumulate wealth to self-insure against shocks to their earnings. To understand the impact of higher-order earnings risk I compare two economies, a canonical economy and a higher-order economy, which only differ in their earnings dynamics. Log-earnings are in both economies the sum of a standard

Gaussian transitory shock and a persistent component. In the canonical economy the persistent component has regular normal innovations and, thus, the earnings distribution does not capture any higher-order moments. In the higher-order economy the persistent component is instead calibrated to match the unconditional earnings distributions from recent empirical studies, in particular Guvenen et al. (2014). The second contribution is methodological. I use a new global solution method to solve for the aggregate dynamics of heterogeneous-agent macro models. While the literature on solution methods for heterogeneous-agent macro models is rapidly growing much of the recent progress pertains to the development of perturbation methods (Auclert et al. (2021), Winberry (2018)) and less to global solution methods such as Fernández-Villaverde et al. (2020) and Schaab (2021).

First, I find that earnings inequality increases over the long run in the presence of higher-order moments. Induced by more severe negative shocks, earnings of the bottom quintiles of the earnings distribution decrease relative to the average and top earners. Moreover, matching moments of the 1 and 5-year earnings growth rate distributions induces shocks to earnings to become more persistent. Income and wealth poor households respond to lower earnings and less upside potential by reducing consumption and increasing savings. In contrast, wealthy households behave as permanent income consumers and therefore barely respond to changes to higher-order moments of shocks to their earnings. Since consumption and wealth are strongly correlated with earnings, the rise in earnings inequality is passed through to larger consumption inequality and lower wealth inequality, the wealth Gini falls from 0.71 to 0.68. To put this into perspective, in order to match the same wealth Gini of 0.71 as in the canonical economy the higher-order economy requires discount heterogeneity to increase by 20%. The mild looking changes in the cross-section mask strong consumption responses by the poor. In particular, consumption as a share of cash at hand increases by roughly 13 percentage points for the bottom quintile of the wealth distribution. Thus, the presence of higher-order earnings risk increases the importance of wealth and, thus, precautionary savings as an insurance tool for the poor. This reinforces the known difficulty of standard heterogeneous-agent models to match the empirically observed dispersion of wealth when accounting for higher-order earnings risk. Instead these findings put more emphasis on alternative sources of wealth inequality such as heterogeneity in asset returns which has recently received much attention. [1]

Second, higher-order earnings risk changes the dynamics of the wealth distribution over the business cycle only for the bottom wealth quintiles, leaving the time series of economic

[1] For example, Fagereng et al. (2020) and Bach et al. (2020) find substantial heterogeneity in returns to wealth, Hubmer et al. (2021) argue in a large-scale heterogeneous-agent model that asset return heterogeneity is key to matching the time series of wealth inequality.

aggregates such as capital and output mostly unchanged. As poor households increase savings to compensate worse and more persistent shocks to their earnings their capital income increases. Moreover, the higher-order moments of earnings, particularly excess kurtosis, are partially passed through to consumption and wealth holdings of the bottom wealth quintile. Both effects reduce the correlation between consumption of the wealth poor and aggregate output substantially, from 0.77 in the canonical economy to 0.19 in the higher-order economy. Lastly, welfare costs of higher-order earnings risk are concentrated at impatient households at the bottom of the wealth distribution who would pay 1.7% of permanent consumption to remain in the canonical economy.

Methodological Contribution. I use General Polynomial Chaos Expansion (GPCE) to solve for the aggregate dynamics of the model. I build on Proehl (2017) and develop a form of GPCE that is suitable as a solution method for heterogeneous-agent macro models and show that it generates a law of motion which accurately forecasts aggregate prices in the presented model. As an extension and outlook I solve a model with time-varying earnings risk and show that the method also performs well in that context.

GPCE is a global projection method that expands the cross-sectional household distribution μ_t in terms of a series of orthogonal polynomials Ψ_i and thereby approximates the distribution with time-varying coefficients $\alpha_{i,t}$.

$$\mu_t(s_t) = \sum_{i=1}^{n} \alpha_{t,i} \Psi_i(\nu) \tag{1.1}$$

where s_t is a vector of individual state variables and ν is a base random variable based on which the polynomials are generated. GPCE gains its efficiency through a smart choice of the base random variable, in particular a base random variable that is distributed similar to the cross-sectional household distribution. This allows GPCE to achieve a high degree of accuracy with a low dimensional approximation and thereby overcome the curse of dimensionality. I make two contributions with respect to making GPCE suitable in the first place and also efficient for heterogeneous-agent macro models:

1. One underlying assumptions when using GPCE is that the model parameters are independent. When approximating the cross-sectional household distribution this assumption requires the individual state variables to be independent which does not hold in most economic settings and, particularly in heterogeneous-agent models, is in stark contrast with the empirical motivation for the models. I develop a GPCE method that allows for dependence between the individual state variables and thereby makes GPCE suitable to solving heterogeneous-agent models.

2. I show that a short outer iteration scheme that updates the base random variable such that it is similarly distributed as the ergodic household distribution leads to reduced and less biased forecast errors for a given number of aggregate state variables. Moreover, I implement a version of the method that projects on different bases depending on whether the economy is in a recessions or an expansions. This is of particular importance when applying GPCE to economies which at times move far away from the ergodic distribution since it allows to choose the appropriate base distribution and therefore polynomials to project on for different regions of the aggregate state space.

The remainder of the paper is structured as follows. The next section places this paper in the literature. Section 3 describes the quantitative model used to analyze the implications of higher-order earnings risk. In section 4 I present the global solution method as well as my contributions in detail. Sections 5 and 6 calibrate the model and present the economic results. Section 7 evaluates the accuracy of the computational method in contrast to previously used ones and its ability to accurately solve models which exhibit stronger aggregate non-linearities than the economy analyzed here. Section 9 provides an outlook for this project, section 10 concludes.

2 Related Literature

This paper is related to several branches of the literature. First and foremast, it relates to the long standing literature studying theories of wealth inequality starting with Bewley (1977), mrohorolu (1989), Huggett (1993), and Aiyagari (1994). The core of the Hugget-Bewely-Aiyagari (BHA) economy builds on idiosyncratic uninsurable shocks to households' earnings. Households have to accumulate non-state-contingent assets to smooth consumption. The shocks endogenously generate ex-post dispersion in wealth as households experience different shocks over time and, thus, have different ability to accumulate wealth. In the basic model wealth inequality is completely determined by the exogenous specification of the earnings process, however, properly calibrated earnings processes fall short of generating the dispersion of wealth observed in the data. An extensive literature builds on the core model and introduces additional sources of wealth inequality such as entrepreneurial ability (Cagetti and De Nardi (2006), Quadrini (2000)) or more recently heterogeneity in returns to wealth (Hubmer et al., 2021). Starting with Krusell and Smith (1998) a large literature has also explored the dynamics of earnings and wealth inequality in the presence of business cycle fluctuations which affect households across the wealth distribution differently. I contribute to this literature by revisiting the role of earnings dynamics for wealth dispersion, in particular focusing on the role of higher-order earnings risk.

This paper is further related to the growing literature studying the implications of higher-order earnings dynamics for consumption and wealth. In particular, a more recent empirical body of papers documents that, in contrast to previous work, individual earnings shocks exhibit sizable higher-order moments in form of negative skewness and excess kurtosis. Incorporating these higher-order dynamics in richer earnings processes has been shown to matter for a variety of economic questions. Golosov et al. (2016) show that it has important implications for optimal redistribution and insurance. In a standard Aiyagari economy Civale et al. (2017) evaluate the implications of higher-order earnings risk for the aggregate capital stock. De Nardi et al. (2018) and Busch and Ludwig (2020) both analyze the implications for consumption-savings decisions over the life-cycle in partial equilibrium life-cycle models. They find that a richer process for earnings increases consumption insurance and moves insurance against persistent shocks closer to the data. Importantly, Busch and Ludwig note that the increased average consumption insurance masks the fact that insurance against negative shocks nevertheless falls. The intuition is that negative skewness coupled with excess kurtosis induces negative shocks to be larger in size and the rise in precautionary savings is not sufficient to fully offset those. I contribute to this literature by analyzing the heterogeneous role of higher-order earnings risk for consumption and wealth across the wealth distribution. However, in contrast to much of the existing literature (De Nardi et al., 2018; Busch and Ludwig, 2020) I allow for general equilibrium effects and business cycle fluctuations. This is important as wealth inequality is not only determined by individual household responses to earnings dynamics but also through their interactions in capital and labor markets which are themselves subject to business cycle fluctuations.

My work is further related to the extensive literature on solution methods for heterogeneous-agent models, and in particular global solution methods [2]. In their seminal contribution Krusell and Smith (1998) develop a global solution method that proposes a parametric law of motion for the aggregate capital stock and is still widely used. More recently, there has been a series of papers on perturbation methods using similar finite-dimensional distribution approximations as I do (Boppart et al. (2018), Winberry (2018)). I build on and my contribution is most closely related to Proehl (2017) who first introduced General Polynomial Chaos Expansion (GPCE) to an economic setting and whose paper provides important insights into how to implement GPCE in practice. Much of her focus is on the theoretical foundations of the method, proving convergence to the rational expectation equilibrium in a standard Krusell-Smith economy. My paper is also closely related to Schaab (2021) who developes a similar global solution method. While he generates the basis functions to project on in a different way, he also focuses on how to choose the basis functions optimally in order to generate a low-order efficient and non-parametric approximation for the law of motion. I contribute to that literature by developing an efficient version of GPCE that, crucially, allows for individual state variables to be dependent, thereby making it a suitable solution method for heterogeneous-agent models.

[2] e.g. Den Haan (1996), Den Haan (1997), Reiter (2010), Maliar et al. (2010), Proehl (2017), Fernández-Villaverde et al. (2020), Schaab (2021)

3 Model

The model builds on Krusell and Smith (1998), thus, it is a general equilibrium model with household heterogeneity and aggregate risk. It differs from Krusell and Smith (1998) at the microeconomic level in two important ways: First, households experience idiosyncratic shocks to their earnings rather than to their employment status. This allows the model to match the estimated earnings dynamics from the data and to obtain more realistic cross-sectional earnings and wealth distributions. Second, households differ in their permanent discount factors which is a common and known modification to match the cross-sectional dispersion in wealth that we observe in the data. Generating a realistic degree of wealth dispersion is necessary to understanding the implication of higher-order earnings risk as households across the wealth distribution respond differently to earnings risk. Moreover, calibrating the dispersion of discount factors to match cross-sectional moments of the wealth distribution gives rise to an intuitive measure of how much higher-order earnings risk affects wealth dispersion, that is, the change in discount factor dispersion required to match the same moments.

3.1 Technology

A representative firm maximizes profits by renting capital K_t and labor L_t from households to produce a non-storable consumption good Y_t. The firm takes factor prices for capital and labor, R_t and W_t, as given and operates a standard Cobb-Douglas technology subject to aggregate productivity shocks z_t. The static maximization problem is given by

$$\max_{K_t, L_t} Y_t - R_t K_t - W_t L_t.$$

where

$$Y = z_t K_t^{\alpha} L_t^{1-\alpha} \tag{3.1}$$

and α denotes the capital share of output. The mean-reverting productivity shock z_t is the source of aggregate uncertainty in this model and follows an AR(1) process given by

$$\log(z_t) = (1-\rho_z)\mu_z + \rho_z \log(z_{t-1}) + \epsilon_t^z,$$

with $\mathrm{E}[z_t] = 1$. Input markets are competitive, thus, factor prices R_t and W_t are equal to their marginal products. Capital used in production depreciates at rate δ.

3.2 Households

The economy is populated by a continuum of households of unit mass. Households survive from each period to the next with constant probability θ as in Blanchard (1985). Carroll et al. (2015) show that this ensures the existence of an ergodic wealth distribution. Each period a mass $1-\theta$ of new households is born with zero initial wealth, thus, leaving the overall population size unchanged.

3.2.1 Preferences

Households maximize expected discounted utility from consumption with standard time-separable preferences, that is, period utility $u(c)$ is continuous, strictly increasing and concave. Households differ in wealth, labor productivity and discount factors. Discount factor heterogeneity as a tool to generate a degree of wealth inequality similar to the data was first introduced by Krusell and Smith (1998) who postulated that households face stochastic shocks to their discount factors. Instead, I follow Carroll et al. (2017) and assume households have different but permanent discount factors. This can be interpreted as capturing a variety of channels of heterogeneity such as differences in risk preferences, age or expectations that matter for consumption-savings decisions, and thereby for the resulting distribution of wealth. Discount factors are distributed uniformly on the intervall $[\beta-\Delta, \beta+\Delta]$ which I discretize with three possible values. Thus, there are two parameters, the mean discount factor β and the dispersion Δ, that need to be chosen in the calibration.

3.2.2 Household Problem

Each period households choose how much to consume and save while supplying labor inelastically to the firm. Households are subject to uninsurable idiosyncratic earnings risk as well as aggregate uncertainty, they can save via risky capital and have otherwise only access to perfect annuity markets. Assets of the deceased are distributed equally among the surviving population and the newborn households are born with zero assets.

Let a_t and y_t denote current asset holdings and labor productivity, respectively, and β_i the permanent discount factor of households of type i. Further, let Z_t be a vector of aggregate state variables consisting of aggregate productivity z_t and the cross-sectional household distribution $\mu_t(\beta_i, a_t, y_t)$ which households need to know in order to predict future prices. Then the recursive household problem is given by

$$V(\beta_i, a_t, y_t, Z_t) = \max_{c_t, a_{t+1}} u(c_t) + \theta \beta_i \mathrm{E}_t \left[V(\beta_i, a_{t+1}, y_{t+1}, Z_{t+1}) \right] \tag{3.2}$$

subject to budget constraint, borrowing constraint and aggregate law of motion

$$\begin{aligned}
&c_t + a_{t+1} \leq a_t(1 + R(Z_t) - \delta) + y_t W(Z_t) \\
&a_{t+1} \geq 0 \\
&\mu_{t+1} = A(Z_t, z_{t+1})
\end{aligned}$$

Household utility is given by

$$u(c_t) = \frac{c_t^{1-\sigma}}{1-\sigma} \tag{3.3}$$

where σ as usual quantifies risk aversion and the inverse of the intertemporal elasticity of substitution.

3.2.3 An Earnings Process with Higher-Order Moments

The focus of this project is the effect of higher-order idiosyncratic earnings risk on household behavior. Labor productivity y_t is the sum of a transitory and a persistent component

$$y_t = \exp\left(p_t + \epsilon_t\right), \qquad \epsilon_t \underset{iid}{\sim} \mathcal{N}(0, \sigma_\epsilon^2) \tag{3.4}$$

The persistent component p_t follows an AR(1) process given by

$$p_t = \rho p_{t-1} + \eta_t, \qquad \eta_t \underset{iid}{\sim} \mathcal{F} \tag{3.5}$$

The distribution of innovations to persistent earnings $\mathcal{F}$ is calibrated to match recent empirical findings from Guvenen et al. (2015) and exhibits negative skewness and excess kurtosis. Let $\Pi(y_{t+1}|y_t)$ denote transition probabilities and $p(y)$ the invariant distribution of overall earnings that come out of the calibration. I propose a simple way to discretizing the persistent component when trying to match higher-order moments of earnings building on Civale et al. (2017) which, in contrast to existing methods, does not require a simulation step [3].

[3] See appendix B for a detailed description of the discretization method.

3.3 Recursive Competitive Equilibrium

A recursive competitive equilibrium consists of policy and value functions for households in all individual and aggregate states, $c_t(\beta_i, a_t, y_t, Z_t)$, $a_{t+1}(\beta_i, a_t, y_t, Z_t)$, $V(\beta_i, a_t, y_t, Z_t)$, functions for aggregate prices, $R(Z_t)$ and $W(Z_t)$, and an aggregate law of motion for the cross-sectional household distribution $A(Z_t, z_{t+1})$ such that

1. Given pricing functions and the aggregate law of motion, $R(Z_t)$, $W(Z_t)$ and $A(Z_t, z_{t+1})$, $V(\beta_i, a_t, y_t, Z_t)$ solves the household problem described in (3.2), and $c_t(\beta_i, a_t, y_t, Z_t)$, $a_{t+1}(\beta_i, a_t, y_t, Z_t)$ are the corresponding decision rules.

2. Factor prices are equal to their marginal products and given by

$$R(Z_t) = z_t \alpha \left(\frac{K_t(Z_t)}{L_t} \right)^{\alpha - 1}$$
$$W(Z_t) = z_t (1 - \alpha) \left(\frac{K_t(Z_t)}{L_t} \right)^{\alpha}$$

3. Capital and Labor markets clear

$$K_t(Z_t) = \int a_t \; d\mu(\beta_i, a_t, y_t)$$
$$L_t = \bar{L} = \sum_y y p(y)$$

4. The aggregate law of motion A used by households to forecast the evolution of the cross-sectional distribution is consistent with the realized law of motion induced by household behavior and the exogenous processes for aggregate and idiosyncratic risk.

4 The Global Solution Method

This section discusses the methodological contribution and presents the global solution method used to solve for the aggregate dynamics of the economy. Reading this section is not required in order to follow the economic results.

It is well known that heterogeneous-agent models with aggregate non-linearities are difficult to solve. The reason is that households base current actions on forecasts of future prices and in any rational expectations equilibrium those forecasts have to be consistent with how the economy actually evolves, that is, with the future prices that materialize. If the optimal policy function for savings features a constant propensity to save out of current assets and income, the evolution of aggregate capital can be described by an average savings function of a representative household and the current capital stocks becomes a sufficient statistic for future prices. However, if the propensity to save out of current assets and income is not constant across the individual state space, forecasting the aggregate capital stock requires forecasting the evolution of the cross-sectional household distribution. The distribution then becomes an aggregate state variable which makes the aggregate state space infinite-dimensional. The goal of any solution method is to accurately approximate the aggregate state space in finitely many dimensions, in fact in as few as possible.

4.1 Polynomial Chaos Expansion

General Polynomial Chaos Expansion (GPCE) is a projection method that represents some random variables of interest, such as the cross-sectional household distribution, as a function of a basic random variable with some specified distribution. The function is an infinite series of orthogonal polynomials which maps the basic random variable into the space of square-integrable random variables. Thus, the basic idea is to represent the cross-sectional household distribution as a random variable by approximating it as a series of polynomials which themselves are random variables.

I build on Proehl (2017) who first used GPCE to solve a standard Krusell-Smith economy. Their focus was largely on the theoretical underpinnings, among other things proving convergence to the rational expectations equilibrium. The central task for any projection method is to choose the projection base, here the set of polynomials. The advantage of polynomial chaos expansion over other projection methods is that when the basic random variable is chosen such that it is similar to cross-sectional household distributions to be approximated the order of polynomials required to obtain a good approximation of the household distribution is significantly reduced. Moreover, given a base random variable the set of polynomials can be generated with standard methods. However, when approximating multivariate distributions GPCE builds on the assumption that the variables that make up the multivariate distribution are independent. When applied to economic models, this assumption requires the individual state variables to be independent. In most economic settings individual state variables are not independent, and in heterogeneous-agent models such as the one presented here generating a realistic joint distribution of income and wealth is instead a central motivation. The standard GPCE method, therefore, cannot accurately approximate the law of motion in these models because the implied independence assumption causes the distribution of the projection base to miss important features of the cross-sectional household distribution to be approximated. To make GPCE suitable for this class of models, I develop an approach to account for dependent base variables and show that it generates an accurate law of motion.

Formally, let $\nu \sim \mathcal{F}$ be some basic random variable and Ψ_i a series of orthogonal polynomials which map ν into the space of square-integrable random variables. Then any square-integrable random variable μ_t can be written as

$$\mu_t = \sum_{i=0}^{\infty} \alpha_{t,i} \Psi_i(\nu) \tag{4.1}$$

where $\alpha_{t,i}$ is a series of scalar coefficients. Here μ_t is a cross-sectional household distribution which is fully characterized by the specific series of coefficients $\{\alpha_{t,i}\}_{i=1}^{n}$. In practice, the above polynomial expansion must be truncated for the aggregate state space to become finite-dimensional and to make the household problem tractable, and thereby computationally feasible. Let the order of truncation be denoted by n, thus, one replaces the true household distribution with an n-th order GPCE approximation. Then the evolution of the approximate household distribution $\mu_{t,n}$ is fully characterized by the evolution of coefficients $\{\alpha_{t,i}\}_{i=1}^{n}$. The number of dimensions in the aggregate state space required to approximate the cross-sectional distribution reduces to n and the household problem becomes tractable. Naturally, truncating the infinite series in equation 4.1 introduces an approximation error. The quality of the solution method is determined by its ability to approximate the distributions of interest sufficiently well while keeping the order of truncation n low and, thus, the aggregate state space small.

4.1.1 Choice of the Base Random Variable

The first task is to choose the distribution of the base random variable which the polynomials take as an argument. It is crucial that this base distribution closely mimics the cross-sectional household distributions that is approximated in the dynamic model. This will allow the truncation order to be low while still achieving a high degree of accuracy. I propose an outer iteration scheme that chooses and adjusts the base distribution in the following way:

1. Begin with an inital base distribution. Here the cross-sectional household distribution from the steady state model without aggregate risk is usually a good choice. First, the steady state distribution is easy to obtain and often one solves for the steady state anyway in order to obtain an initial guess for the policy functions. Second, the shape of the ergodic distribution from the dynamic model is often similar to the shape of the steady state distribution.

2. Solve the model with the polynomials based on the initial base distribution and simulate the economy. Compute forecast errors and the ergodic distribution.

3. If forecast errors are too large, update the base distribution by using the ergodic distribution from the simulation. In practice it is useful to dampen the adjustment of the base distribution by using a convex combination of the ergodic distribution and previous base distribution. This is repeated until the forecast errors of the law of motion are sufficiently small or the ergodic and base distributions are equivalent, thus, making further updates redundant. If that is the case while the forecast errors are still too large the order of truncation needs to be increased and the scheme repeated.

The overall shape of the ergodic distribution is usually not sensitive to small forecast errors in the law of motion, thus, only a small number of iterations is necessary until the base and ergodic distribution converge, for example I used three iterations.

The implementation as described above may not work well in economic settings in which parts of the aggregate state space exhibit strong non-linearities. In that case the cross-sectional household distribution can move far away from the ergodic distribution in those parts of the aggregate state space. This is often the case in models with financial accelerators such as Brunnermeier and Sannikov (2014) or Fernández-Villaverde et al. (2020). In that case, one can project on different base distributions and, thus, different sets of polynomials depending on where in the aggregate state space the economy is. For example, in models with financial accelerators the level of leverage of the financial sector could determine the current base distribution. In section 8 I solve an economy with cyclical earnings risk by projecting on different bases in recessions and expansions, thus, the current base distribution depends on the aggregate level of productivity z_t. The method provides an accurate law of motion in that setting as well.

4.1.2 Polynomials with Standard GPCE

Given a choice for the base random variable ν one needs to generate the set of orthogonal polynomials $\{\Psi_i(\nu)\}_{i=1}^n$ on which to project. When approximating multivariate distributions, such as the cross-sectional household distribution over (β, y, a) in the presented model, the standard GPCE method prescribes to choose an independent basic random variable $(\nu_1, \nu_2, \nu_3) = (\nu_\beta, \nu_y, \nu_a)$ for each dimension. For each base random variable ν_j the respective univariate set of polynomials $\{\Psi_{j,i}(\nu_j)\}_{i=1}^n$ is generated separately. The polynomials can be generated using the three-term recurrence relation (see e.g. Gautschi (1982) or Zheng et al. (2015)) [4]. The multivariate set of orthogonal polynomials $\{\Psi_i(\nu)\}_{i=1}^n$ is then a tensor product of the univariate polynomials given by

$$\begin{aligned}
\mu_t &= \sum_{i=1}^{\infty} \alpha_{t,i} \Psi_i(\nu) \\
&= \sum_{i=1}^{\infty} \alpha_{t,i} \Psi_i(\nu_\beta, \nu_y, \nu_k) \\
&= \sum_{i=1}^{\infty} \alpha_{t,i} \sum_{i_1, i_2, i_3 = 1}^{i} \Psi_{1,i}(\nu_\beta) \cdot \Psi_{2,i}(\nu_y) \cdot \Psi_{3,i}(\nu_a) \; \mathbb{1}\{i_1 + i_2 + i_3 = i\} \qquad (4.2)
\end{aligned}$$

The corresponding probability density function $f(\nu)$ of the base distribution ν is also given by the tensor product of the univariate pdf's, that is, the joint distribution of the base random variable $\nu = (\nu_\beta, \nu_y, \nu_a)$ is generated by multiplying the marginal pdf's. However, in the presented model the individual state variables (β, y, a) are not independent but instead exhibit strong correlation. Generating the set of polynomials under the independence assumption as described above generates a base distribution with a set of corresponding polynomials that does not closely mimic the cross-sectional household distribution to be approximated. This leads to large approximating errors in the law of motion if accounting for the correlation between the individual state variables is relevant for the evolution of prices. For example, this is the case if asset poor agents behave significantly different depending on whether their earnings are low or high or depending on whether they are impatient or patient. As shown and discussed in section 7 this applies to my model.

[4] See appendix A for details on generating the polynomials with the the three-term recurrence relation

4.1.3 A Dependent GPCE Method

To allow for dependence among the base random variables, I instead propose a different way of generating the polynomials and using GPCE. Instead of approximating the joint, unconditional cross-sectional household distribution with a joint base random variable and joint set of polynomials as described in the previous section, I approximate the evolution of the conditional asset distribution $\mu_t(a|\beta, y)$, thus for each (β, y) group, separately. I choose a univariate base random variable $\nu(a|\beta, y)$ and generate a set of univariate polynomials $\{\Psi_i(\nu(a|\beta, y))\}_{i=1}^n$ for each (β, y). As a result, I effectively apply GPCE to each (β, y) asset distribution separately and, thus, the evolution of each (β, y) asset distribution is described by a set of coefficients $\{\alpha_{t,i}(\beta, y)\}_{i=1}^n$. This naturally allows the different base random variable to depend on the discount factor β as well as income y, and thereby allows individual state variables to be dependent. Solving for a law of motion for all $\{\alpha_{t,i}(\beta, y)\}_{i=1}^n$, however, would add n aggregate state variables for each (β, y) group. In order to overcome the resulting curse of dimensionality I solve for the law of motion only at specific combinations of coefficients in the aggregate state space.

I illustrate this by looking at the first-order coefficients $\alpha_{t,1}(\beta, y)$. The first-order polynomials are by default equal to 1, $\Psi_i(\nu) = 1$ and higher-order order polynomials are constructed with zero-mean thereby making the first-order coefficients $\alpha_{t,1}$ equal to the mean of the distribution at time t. First, I discretize the state space of coefficients $\alpha_{t,1}(\beta, y)$ separately for each (β, y) base using the same number of grid points for each one. This yields the first-order coefficients to have different grids across (β, y) groups. Therefore, it allows exactly defining the possible average asset holdings for each (β, y) conditional asset distribution at which the law of motion is solved for. This is not possible with the standard GPCE approach.

In order to overcome the curse of dimensionality, I only solve for points in the aggregate state space where all $\alpha_{t,1}(\beta, y)$ have the same index in their respective grids. For example, if each grid for the first-order coefficients $\alpha_{t,1}(\beta, y)$ has a total of two grid points I only solve for the law of motion at points in the aggregate state space where all $\alpha_{t,1}(\beta, y)$ are either low or high. Economically, I thereby assume that the average asset holdings of (β, y) groups are positively correlated. For the sake of forecasting prices it is therefore sufficient to approximate the law of motion at points in the aggregate state space where average asset holdings for all conditional asset distributions $\mu_t(a|\beta, y)$ are simultaneously low or high. Importantly, this assumption does not imply that average asset holdings by (β, y) groups simultaneously increase and decrease by the same amount or even factor but rather that the sign of changes in the average asset holdings is the same over time. The aggregate law of motion then takes as input not the first-order coefficient for all (β, y) asset distributions but instead I choose the weighted average across all groups resulting in an unconditional coefficient

$$\alpha_{t,1} = \sum_{\beta,y} \alpha_{t,1}(\beta, y)p(\beta, y) \tag{4.3}$$

The weights $p(\beta, y)$ are given by the exogenous, invariant mass of agents with individual state (β, y). As a result, the first-order coefficient $\alpha_{t,1}$ will be equal to the aggregate capital in the economy as with standard GPCE despite generating each (β, y) asset distribution with their own separate set of polynomials and coefficients.
I then apply the same assumption when generating higher-order polynomials and discretizing the state space of the corresponding coefficients. As a result, the total number of aggregate grid points used to approximate the cross-sectional distribution is again equal to the order of truncation n.

4.1.4 Simultaneously Solving for the Transition and Policy Functions

I will briefly outline the specific steps taken when solving for the law of motion for the unconditional coefficients $f(\{\alpha_{t,i}\}_{i=1}^n) = \{\alpha_{t+1,i}\}_{i=1}^n$.
The traditional way of finding a law of motion for the cross-sectional household distribution that is consistent with household behavior is to start with a parametric guess for the law of motion and then solve and simulate the model. The guessed law of motion is updated based on the realized law of motion in the simulation and the model is solved again with the updated law of motion. It usually requires many rounds of simulation until the guessed and simulated laws of motion converge. A major advantage of projection methods such as Dependent GPCE as presented here or the one presented in Schaab (2021) is that they do not need the simulation step but instead solve for a nonparametric law of motion while solving for the policy functions.

First, one initializes the algorithm with some guess for the transition function just as one begins with some initial guess for the policy functions. In this class of models a good initial guess is usually to assume that the cross-sectional household distribution does not change. Thus, given a set of coefficients $\{\alpha_{t,i}\}_{i=1}^n$ in the discretized aggregate state space the initial guess for the law of motion at that point is given by

$$\begin{aligned} f(\{\alpha_{t,i}\}_{i=1}^n) &= \{\alpha_{t+1,i}\}_{i=1}^n \\ &= \{\alpha_{t,i}\}_{i=1}^n \end{aligned}$$

In a given aggregate state $(z_t, \{\tilde{\alpha}_{t,i}\}_{i=1}^n)$ the law of motion and current policy functions determine the possible aggregate and individual states next period and, thus, the right-hand-side of the Euler equation. One then solves the Euler equation and updates the current guess for policy functions in all individual states. The updated policy functions are then used to update the law of motion in the following way. The current conditional asset distributions $\tilde{\mu}_t(a|\beta, y)$ for all (β, y) is given by the polynomial expansion in the current conditional coefficients

$$\tilde{\mu}_t(a|\beta, y) = \sum_{i=1}^{n} \tilde{\alpha}_{t,i}(\beta, y)\Psi_i(\nu(a|\beta, y))$$

Given the policy function for asset holdings $a_{t+1}(z_t, \{\tilde{\alpha}_{t,i}\}_{i=1}^n, \beta, y, a)$, the conditional asset holdings next period $\tilde{\mu}_{t+1}(\bar{a}|\bar{\beta}, \bar{y})$ at some specific point $(\bar{a}, \bar{\beta}, \bar{y})$ in the individual state space evolves according to

$$\tilde{\mu}_{t+1}(\bar{a}|\bar{\beta}, \bar{y}) = \int_{\beta,y,a} p(a|\beta, y)\mathbb{1}\{a_{t+1}(z_t, \{\tilde{\alpha}_{t,i}\}_{i=1}^n, \beta, y, \tilde{\mu}_t(a|\beta, y)) = \bar{a}\}p(\bar{\beta})p(\bar{y}|y)d(\beta, y, a)$$

where $p(\bar{\beta})$ is the exogenous, invariant mass of agents with $\beta = \bar{\beta}$, $p(\bar{y}|y)$ is the exogenously determined transition probability for earnings, and $p(a|\beta, y)$ is the conditional base pdf of group (β, y). One then projects the realized cross-sectional household distribution $\tilde{\mu}_{t+1}$ onto each polynomial in order to obtain the realized conditional coefficients next period.

$$\begin{aligned}
\alpha_{t+1,i}(\beta, y) &= \frac{\langle \tilde{\mu}_{t+1}, \Psi_i \rangle}{\langle \Psi_i, \Psi_i \rangle} \\
&= \frac{1}{\langle \Psi_i, \Psi_i \rangle} \int_{\beta,y,a} \tilde{\mu}_{t+1}(a|\beta, y)\Psi_i(a|\beta, y)dF(a|\beta, y) \qquad \forall i
\end{aligned}$$

where $dF(a|\beta, y)$ denotes the conditional density of the base random variable for group (β, y). The transition function is then updated with the new unconditional coefficients $\alpha_{t+1,i}$ which follow from 4.3. This is repeated until policy functions and transition function converge.

5 Taking the Canonical Economy to the Data

The model is calibrated to match US data at an annual frequency. I calibrate the model with the canonical specification and then use the same set of parameters in the economy with higher-order earnings risk. The distribution of discount factors is calibrated internally, the remaining parameters are calibrated externally.

Aggregate Risk and Technology. As standard in the literature, I model aggregate shocks as a two-state AR(1) process where the two states correspond to recessions and expansions $z \in \{z_r, z_e\}$. Following Krusell and Smith (1998) and much of the literature I set the standard deviation of aggregate productivity to $\sigma_z = 0.01$ to roughly match the magnitude of business cycle fluctuations in the US. This results in productivities in recessions and expansions to be $z_r = 0.99$ and $z_e = 1.01$, respectively. I follow Busch and Ludwig (2020) in their classification of recessions and expansions which gives rise to the following transition matrix for aggregate productivity

$$\pi(z_{t+1}|z_t) = \begin{pmatrix} 0.388 & 0.612 \\ 0.231 & 0.769 \end{pmatrix} \tag{5.1}$$

and corresponding invariant distribution $\pi_z = (0.274, 0.726)$. The capital share is set to $\alpha = 0.36$ and capital depreciates at rate $\delta = 0.08$. Average income $\bar{y}$ in the economy is normalized such that average output over the business cycle is equal to 1.

Preferences and Discount Factor Distribution. Households are assumed to have logarithmic utility. I calibrate the parameters of the discount factor distribution, meaning the average discount factor β and the dispersion Δ, to match an average capital to output ratio of 3 and a Gini coefficient for wealth of 0.71. I choose a wealth Gini at the lower end of estimates for the US to prevent the degree of discount factor heterogeneity from becoming too large as it is the only source of wealth inequality in this model beyond earnings inequality. I set the survival probability ϕ to 0.9833 such that the average life-time of an agent is 60. Table 1 provides a summary of the model parameters in the economy with the canonical earnings process.

	Parameter	Value	Source/Target
	Preferences		
σ	Relative Risk Aversion	1	standard
$\bar{\beta}$	Discount Factor	0.9559	K/Y=3
Δ_β	Discount Factor Dispersion	0.0203	Wealth Gini=0.71
ϕ	Survival Probability	0.9833	Avg life-time of 60
	Technology and Aggregate Uncertainty		
α	Capital Share	0.36	standard
δ	Depreciation	0.08	standard
σ_z	Standard Deviation of TFP	0.01	standard
$\bar{y}$	Average Income		E(Y)=1

Table 1: Benchmark Calibration

5.1 The Canonical Earnings Process

The earnings process in the canonical economy does not exhibit any higher-order order moments, the transitory as well as the persistent innovations are Gaussian. Table 2 provides the chosen parameters. There are three parameters to calibrate, the variances of both shocks and the persistence of the persistent component. The objective of this paper is to analyze the impact of higher-order earnings risk by comparing its predictions in a standard heterogeneous-agent model to the canonical earnings process used in much of the literature. Thus, to make my results comparable to similar papers in the literature I set all parameters to standard values that are widely used [5].

Parameter	Value	Description
ρ_η	0.965	Persistence persistent shocks
σ^2_η	0.035	Variance persistent shocks
σ^2_ϵ	0.07	Variance transitory shocks

Table 2: Benchmark Earnings Process

5.2 An Earnings Process with Higher-Order Moments

In order to generate higher-order moments in persistent earnings I let the innovations to the persistent component follow a mixture of normal distributions. The full earnings process is

[5] For example, Krueger and Ludwig (2016), Busch and Ludwig (2020), Storesletten et al. (2004).

then given by

$$y_t = \exp\left(p_t + \epsilon_t\right), \qquad \epsilon_t \underset{iid}{\sim} \mathcal{N}(0, \sigma_\epsilon^2) \tag{5.2}$$

$$p_t = \rho p_{t-1} + \eta_t, \qquad \eta_t \underset{iid}{\sim} \begin{cases} \mathcal{N}(\mu_1, \sigma_1^2) \text{ with } p_1 \\ \mathcal{N}(\mu_2, \sigma_2^2) \text{ with } 1 - p_1 \end{cases} \tag{5.3}$$

I calibrate and discretize the above process with GMM by simultaneously matching a set of moments of the earnings process as well as a set of statistics of the earnings distribution (see appendix B for details). The latter is possible because labor is supplied inelastically, thus, one can compute statistics of the earnings distribution directly from the earnings grid and stationary earnings distribution. For the earnings process I target estimates of earnings moments from Guvenen et al. (2015). While some related paper such as De Nardi et al. (2018) provide discretized processes and transition matrices online there are two advantages from discretizing myself. First, De Nardi et al. (2018) solve a partial-equilibrium model which allows them to use a large number of grid points to obtain a better match, too large to be feasible in my general equilibrium setting with aggregate risk. Second, the provided earnings process exhibits significantly less earnings inequality than observed in the data, however, roughly matching the empirically observed earnings distribution is key when analyzing implications for wealth inequality. Table 3 provides the calibration results for the moments of the earnings process. Apart from moments for 1 and 5-year earnings growth rates from Guvenen et al. (2015), I also target the variance of log earnings from the canonical process to ensure better comparability between the economies.

	Targets	Discretized process
Variance(x)	0.58	0.59
Variance($\Delta_1 x$)	0.23	0.19
Skewness($\Delta_1 x$)	-1.35	-1.39
Kurtosis($\Delta_1 x$)	17.8	15.72
Variance($\Delta_5 x$)	0.46	0.51
Skewness($\Delta_5 x$)	-1.01	-1.25
Kurtosis($\Delta_5 x$)	11.55	9.88

Table 3: Unconditional moments of log earnings for higher-order earnings process, targeted moments from Guvenen et al. (2015)

6 The Cross-Sectional Household Distribution

I first compare the cross-sectional distributions of income, earnings and wealth across the two economies. As a benchmark, I also report the observed cross-sectional facts from the data as estimated by Krueger et al. (2016) and Chang et al. (2019). Matching these cross-sectional distributions is important in a variety of settings, particularly when moving to normative questions such as the welfare costs of business cycles in the presence of cyclical, higher-order earnings risk, and policy design implications. Table 3 provides the ergodic cross-sectional distributions of wealth, earnings and consumption by their respective quintiles from simulating the economies.

	Wealth			**Earn**			**Cons**		
	H-O	Can	Data	H-O	Can	Data	H-O	Can	Data
Bottom 20%	1.6%	1.3%	-0.5%	4.2%	5.8%	4.5%	5.7%	6.8%	6.5%
20%-40%	3.6%	3.1%	0.5%	7.5%	9.7%	9.9%	8.2%	12%	11.4%
40%-60%	6.6 %	5.5%	5.1%	16%	15%	15.3%	18%	16%	16.4%
60%-80%	14%	13%	18.7%	27%	24%	22.8%	28%	24%	23.3%
80%-100%	74%	77%	76.2%	45%	46%	47.5%	39%	42%	42.4%
Gini	0.68	0.71	0.71	0.43	0.41	0.42	0.37	0.35	0.36

Table 4: Cross-sectional wealth, earnings and consumption distributions by quintiles for higher-order (H-O), canonical (Can) and Data.
Data estimates from Krueger et al. (2016) and Chang et al. (2019)

To begin with, the canonical model is able to match the cross-sectional distributions well while only targeting the wealth Gini in the calibration. However, it requires the discount factor dispersion to be relatively large at $\Delta_\beta = 0.0203$ as shown in table 1. Specifically the earnings and consumption distributions are close the data. As for the wealth distribution the model fits the overall shape, however, has the known issue that the bottom quintiles accumulate too much wealth. There are several reasons for that. While the earnings process is calibrated to match earnings dynamics post government transfer it still misses out on some important parts of the social security net for low-income households which reduce their precautionary savings motive. For example, in the canonical economy households in the bottom earnings state earn roughly 8.2% of average earnings, significantly less than what unemployed households usually receive through unemployment insurance in the US. As a comparison, the replacement rate of unemployment insurance for an individual who earned the average wage prior to job loss is 54% according to the OECD. As a result, households at the bottom of the wealth distribution exhibit too strong of a precautionary savings motive. In addition, households cannot borrow in this setting and, thus, the model cannot generate the negative wealth holdings of the bottom quintile observed in the data.

With the same calibration, the model with higher-order earnings risk matches the three cross-sectional distributions still fairly well but worse than the canonical model across the board. The induced earnings distribution is more unequal than in the canonical case and the data. In particular, the bottom quintiles now receive a smaller share of overall earnings. Excess kurtosis and negative skewness increase the dispersion of earnings at the bottom with earnings in the bottom earnings state falling relative to average and top earnings. As a result, more earnings are shifted from the bottom two quintiles to the top three quintiles. In addition to those changes in the earnings distribution the higher-order earnings process is also more persistent than the canonical one. This is mostly induced by matching moments of 1 and 5-year earnings growth rates. As a result, low-income households not only earn less but also are less likely to become lucky and climb up the earnings latter. Since consumption and wealth are both strongly correlated with earnings, the increase in earnings inequality is partly passed through to consumption and wealth. Low-earnings households respond to the change in their earnings prospects by increasing savings and reducing consumption. As a result, consumption of the bottom two quintiles falls while consumption of the top three quintiles rises, leading to an increase in consumption inequality as measured by the Gini coefficient. Similarly, wealth dispersion falls as the bottom quintiles increase their savings which results in a decrease of the wealth Gini from 0.71 to 0.68.

This is relevant for the existing large empirical and quantitative literature that tries to identify potential mechanisms to generate the degree of observed wealth dispersion. These findings suggest that the correlation between earnings and wealth dispersion is weakened by the presence of higher-order risk which in turn increases the importance of alternative sources of wealth dispersion. In order to further quantify the reduction in wealth dispersion I recalibrate the higher-order economy to match the wealth Gini. The required dispersion in discount factors rises by 20% to roughly $\Delta_\beta = 0.025$.

Table 5 reports the earnings distribution by wealth quintiles for the two economies as well as from the data. In line with the analysis above, the share of earnings at the bottom of the distribution falls significantly from 13% in the canonical economy to 10% in the higher-order economy. As a result, the higher-order earnings dynamics help to match the bottom of the distribution as the canonical economy generates too little earnings dispersion among the wealth poor.

	Higher-Order	Canonical	Data
Bottom 20%	10%	13%	9.8%
20%-40%	14%	14%	13%
40%-60%	19%	19%	18%
60%-80%	27%	24%	22%
80%-100%	30%	30%	37%

Table 5: Earnings distribution by wealth quintiles, averaged over all simulation periods. Data estimates from Krueger et al. (2016) and Chang et al. (2019)

In conclusion, the presence of higher-order earnings risk has quantitatively important implications for the bottom of the earnings and consumption distribution while the wealth distribution is only affected mildly. This is mostly due to the fact that the bottom quintiles of the wealth distribution hold little wealth such that even relatively large changes in their consumption-savings behavior do not translate strongly into changes of wealth dispersion.

6.1 Consumption Responses across the Distribution

To get a better understanding of the heterogeneity in responses to higher-order earnings risk that the above described changes in the cross-sectional household distribution imply, a closer look at consumption-savings decisions across the wealth distribution is useful. Table 6 reports consumption as a fraction of cash at hand by wealth quintiles averaged across the respective quintile.

	Higher-Order	Canonical
Bottom 20%	54%	67%
20%-40%	44%	52%
40%-60%	37%	45%
60%-80%	30%	29%
80%-100%	13%	13%

Table 6: Consumption share by wealth quintiles measured as consumption over cash at hand, averaged over all simulation periods.

The table shows that the mild changes in the cross-sectional distributions mask more drastic changes in the behavior of wealth-poor households. The bottom three wealth quintiles reduce their consumption share significantly, with the share falling for the bottom quintile by 13 percentage points. In contrast, the top two wealth quintiles spend roughly the same fraction of asset holdings on consumption as in the canonical economy. Wealthy households behave as permanent income consumers and therefore react little to changes in the higher-order moments of their earnings dynamics but instead primarily care about expected earnings. In order to generate negative skewness while keeping the other moments constant, in particular the variance, additional probability mass has to be moved to high earnings states making good shocks more likely. As a result, expected earnings increase slightly and rich households respond to the increase in permanent income by reducing savings slightly. While the response by the wealthy is very weak it still offsets the increase in savings by the bottom quintiles. The reason is that the top wealth quintile holds roughly 75% of the total wealth in the economy. Aggregate capital and, thus, prices are therefore the same in both economies in equilibrium.

6.2 Time Series Dynamics

The time series dynamics for economic aggregates are very similar in the two models. This is consistent with the above analysis showing that the presence of higher-order earnings risk primarily affects the lowest wealth quintiles while wealthier households behave as permanent income consumers and barely alter their consumption-savings behavior. The lowest wealth quintile holds only a very small share of aggregate wealth and, as a result, changes in its consumption-savings behavior are not passed through to aggregate capital, output and prices. Nevertheless, as consumption behavior of the wealth poor is altered by higher-order earnings risk so are time series dynamics of their consumption. Table 7 shows by wealth quintile the correlation between consumption and aggregate output as well as kurtosis of total wealth held by that quintile.

	corr(C,Y)		**Kurtosis Wealth**	
	H-O	Can	H-O	Can
Bottom 20%	0.19	0.77	6.2	2.4
20%-40%	0.37	0.34	3.5	2.7
40%-60%	0.34	0.42	2.8	3
60%-80%	0.51	0.54	2.5	2.6
80%-100%	0.66	0.61	2.6	2.6

Table 7: Correlation between aggregate consumption of each wealth quintile and aggregate output as well as kurtosis of total wealth held by each quintile.

Higher savings by the poor increase their ability to smooth consumption and, as a result, correlation between aggregate consumption of the bottom wealth quintile and output drops. In addition, the time series of total wealth held by the bottom wealth quintile exhibits substantial excess kurtosis in the higher-order economy. The reason is that for those households cash at hand consists primarily of earnings and, therefore, wealth holdings adopt the higher-order moments of earnings. As a result, variation in consumption of wealth poor households over time is predominantly driven by large changes to their earnings rather than business cycle fluctuations in the higher-order economy. This further reduces the correlation between their consumption and aggregate output compared to the canonical economy.

6.3 Welfare Costs of Higher-Order Earnings Risk

Given the substantial heterogeneity in consumption responses across the wealth distribution it is natural to quantify the cost of higher-order earnings risk in utility terms. In order to do so, I ask in which world households prefer to live. In particular, I look at average expected life-time utility by wealth quintile for each discount factor group and calculate the consumption equivalent variation (CEV) that a household needs to receive in the canonical economy in order to be indifferent between staying in the canonical economy and moving to the higher-order economy. Thus, I assume that households keep their discount factor type and remain in the same wealth quintile among households with the same discount factor. Calculating the CEV separately for each discount factor group is useful as consumption patterns and welfare costs substantially differ between those groups. The CEV for wealth quintile Q of discount factor group β is then given by

$$W(Q,\beta) = \exp\{(1-\beta)\left(V_{\mathrm{HoM}}(Q,\beta) - V_{\mathrm{Can}}(Q,\beta)\right)\} - 1 \tag{6.1}$$

where $V_{\mathrm{i}}(Q,\beta)$ is the average life-time value of households in wealth quintile Q of discount factor group β, averaged across all simulation periods. Table 8 reports those CEV for the patient and impatient households. Note, the wealth quintiles here refer to the wealth distribution of impatient and patient households separately, not the overall wealth distribution.

CEV	Impatient	Patient
Bottom 20%	-1.7%	0.68%
20%-40%	-1.1%	0.42%
40%-60%	1.3%	2.1%
60%-80%	1.9%	2.4%
80%-100%	1.4%	2.2%

Table 8: Average, ergodic CEV for impatient and patient households by wealth quintile of the impatient and patient distribution, respectively.

Consistent with the previous analysis welfare costs of higher-order earnings risk decrease in wealth. An impatient household in the bottom wealth quintile would accept a permanent decrease in consumption of 1.7% in the canonical economy instead of moving to the higher-order economy while remaining impatient and in the bottom wealth quintile of the impatient households. In contrast, wealthy impatient households prefer the higher-order economy and require an increase of 1.3%, 1.9%, and 1.4% in permanent consumption for the top three wealth quintiles respectively in order to stay in the canonical economy. As shown in table 5 the earnings share held by the fourth wealth quintile increases significantly when allowing for higher-order earnings dynamics which is the reason that those household benefit most from the presence of higher-order earnings risk. The overall pattern is similar for patient households though they prefer to live in the higher-order economy across all wealth quintiles. Recall, that these are wealth quintiles of the patient wealth distribution. Since poor patient households respond stronger to higher-order earnings risk and accumulate significantly more wealth than impatient household they are better insured against the additional earnings risk. Consistent with that, the share of patient households in the bottom quintile of the overall wealth distribution falls from 12% in the canonical economy to 8% in the higher-order economy. Wealthy households prefer the higher-order economy as they behave as permanent income consumers and permanent income increases slightly as discussed. Importantly, this analysis abstracts from potential transition dynamics between economies with different earnings dynamics and merely asks which ergodic world households prefer assuming they keep their characteristics. It is therefore intuitive that primarily poor impatient households experience significant welfare costs of higher-order earnings risk because those households stay poor in the higher-order economy.

7 Evaluating the Solution Method

As discussed in 4, I make two contributions that make GPCE not only applicable to heterogeneous-agent models in the first place but also efficient.
First, my approach allows for dependent random base variables and, thus, dependent individual state variables. Second, I implement an outer iteration scheme that adjusts the base random variable, and thereby generates a projection base that more closely mimics the realized distributions in the simulation. In order to understand the corresponding gains I solve the canonical model with three different methods: The here proposed full Dependent GPCE method with both adjustments (full), the Dependent GPCE method without adjusting the initially chosen steady state distribution as base distribution (noBDadj), and finally with the standard GPCE method as used in Proehl (2017) which builds on the independence assumption. In all cases, the order of truncating is three. While forecast errors reduce further with an order of four the gains are relatively small and have no meaningful effect on economic aggregates. I use the canonical economy for this evaluation as the performance of the solution method in the two economies is essentially equivalent.

7.1 One-Period Ahead Forecast Errors

Figure 1 shows boxplots for the distributions of one-period ahead forecast errors for aggregate capital in all three models. Forecast errors are computed as the percentage deviation of forecasted from realized aggregate capital and generated by simulating the economies for 1000 periods (after burn-in).

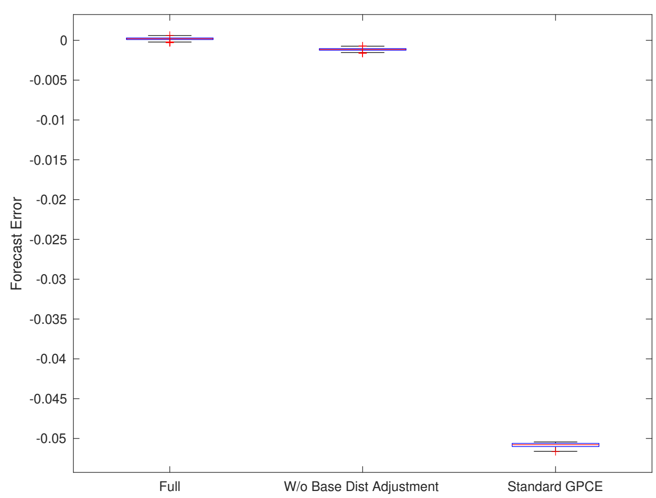

Figure 1: Boxplot for one-period ahead forecasting error distribution, different GPCE approaches

The largest improvement for the accuracy of the law of motion measured at the one-period horizon comes from allowing for the base random variable to be dependent. The average forecast error falls by roughly 3 orders of magnitude, from -0.053 to 0.0002, when moving from the standard to the full Dependent GPCE method. Moreover, the standard GPCE method produces significant bias in its forecasts, systematically forecasting smaller than realized capital stocks. The independence assumption of the standard GPCE method is violated in this model due to correlation between wealth and income as well as wealth and permanent discount factor type. Particularly the latter is strong and therefore responsible for the large forecast errors generated by standard GPCE. The standard Krusell-Smith economy as solved by Proehl (2017) does not exhibit any preference heterogeneity and idiosyncratic income risk only exists in the form of unemployment risk. The unemployed make up a small fraction of the total mass of households and are sufficiently wealthy to have similar propensities to save out of cash at hand as the employed. As a result, the evolution of the economy is almost entirely determined by the employed agents and there is no relevant heterogeneity in individual states beyond wealth dispersion among the employed. However, that is not the case in my model which results in large forecast errors with the standard GPCE method.

To get a clearer picture of the gains from the second contribution figures 2 and 3 plot the forecast error distributions for the Dependent GPCE method with and without base distribution adjustment.

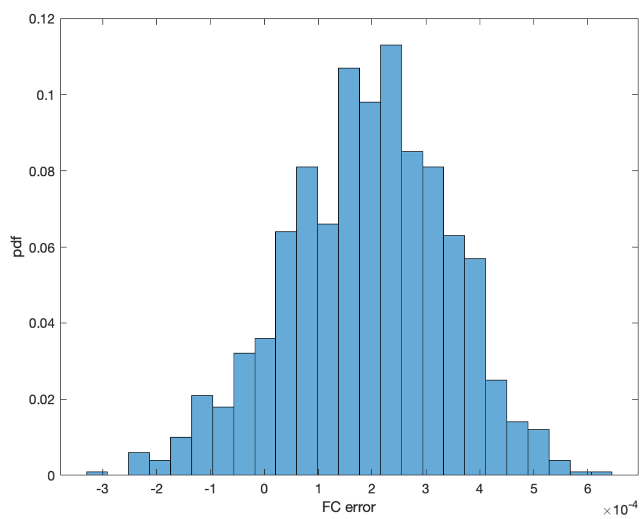

Figure 2: Full Dependent GPCE method

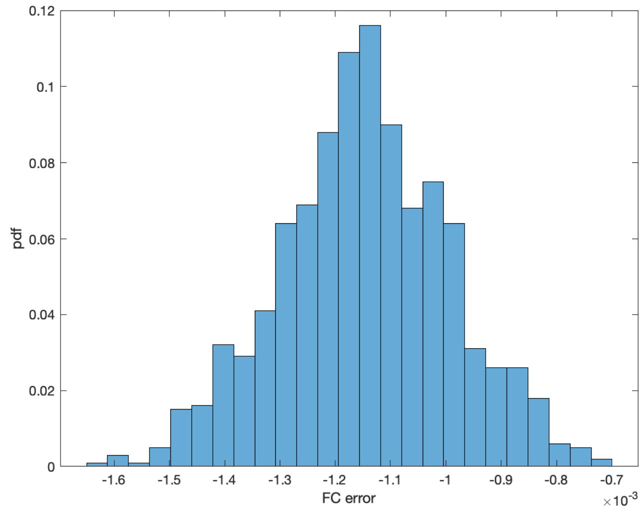

Figure 3: Dependent GPCE method without base distribution adjustment

Adjusting the base distribution to be distributed as the realized ergodic distribution yields two important improvements. First, average forecast errors in absolute terms fall roughly by one order of magnitude. Second, adjusting the base random variable reduces the bias in forecasts. Recall that the base random variable is the basis for the polynomials which in turn determine for a given set of coefficients the sample cross-sectional distributions on the aggregate grid based on which the approximate law of motion is derived. The significant and consistent downward bias in forecasts every period without adjustment shows that the shape of the steady state distribution differs in ways from the realized distributions in the dynamic model that are relevant for the evolution of prices. This is confirmed when comparing specific (β, y) base distributions from the two methods. Figure 4 plots those distributions for the lowest and highest impatient earnings groups for the Dependent GPCE method with (full) and without base distribution adjustment (noBDadj).

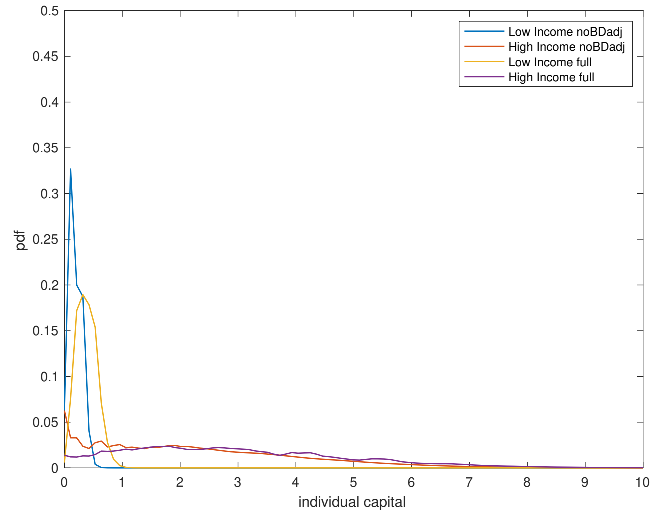

Figure 4: Comparison of base distributions for impatient low earnings and impatient high earnings group across methods

The yellow and blue graphs show the base distributions for impatient households in the lowest earnings group for the Dependent GPCE method with and without base distribution adjustment, respectively. In the dynamic model, income-poor households at the bottom of the wealth distribution have a stronger precautionary savings motive than in the stationary economy due to the presence of aggregate risk and, therefore, accumulate more wealth. Consequently, adjusting the base distribution based on the ergodic distribution significantly reduces the mass of households at or close to the borrowing constraint, resulting in the shift from blue to yellow. The same holds true for the base distributions of high earnings households. Moreover, the shape of the base distributions for low and high earnings households are very different as expected when allowing for dependent base variables because earnings are persistent and, thus, correlated with wealth. In contrast, with the standard GPCE method that relies on the independence assumption the conditional base distributions for different earnings groups are equivalent in their shape by construction.

7.2 Do Forecast Errors Accumulate?

As discussed by Den Haan (2010) one-period ahead forecast errors do not properly assess whether forecast errors accumulate over time. The reason is that current aggregate state variables based on which forecasts are being made are updated every period during the simulation and are, thus, equal to to the true, realized aggregate state variables. In order to evaluate the accuracy of the law of motion further, I therefore check the accumulation of forecast errors. Following Den Haan (2010), one compares the realized aggregate capital stocks from the full simulation to the series of aggregate capital stocks generated by purely iterating over the law of motion. I start with the same initial aggregate state, number of burn-in and simulation periods as in the proper simulation and generate the out-of-sample sequence of aggregate states in the following way

$$
\begin{aligned}
&\{\{\alpha_{0,i}\}_{i=1}^{n}\ , \{\alpha_{1,i}\}_{i=1}^{n}\ , \{\alpha_{2,i}\}_{i=1}^{n}\ , \{\alpha_{3,i}\}_{i=1}^{n}, ...\} = \\
&\{\{\alpha_{0,i}\}_{i=1}^{n}\ , f(\{\alpha_{0,i}\}_{i=1}^{n})\ , f(f(\{\alpha_{0,i}\}_{i=1}^{n}))\ , f(f(f(\{\alpha_{0,i}\}_{i=1}^{n}))), ...\}
\end{aligned}
$$

Figure 5 plots these out-of-sample forecast error distribution for the dependent GPCE method.

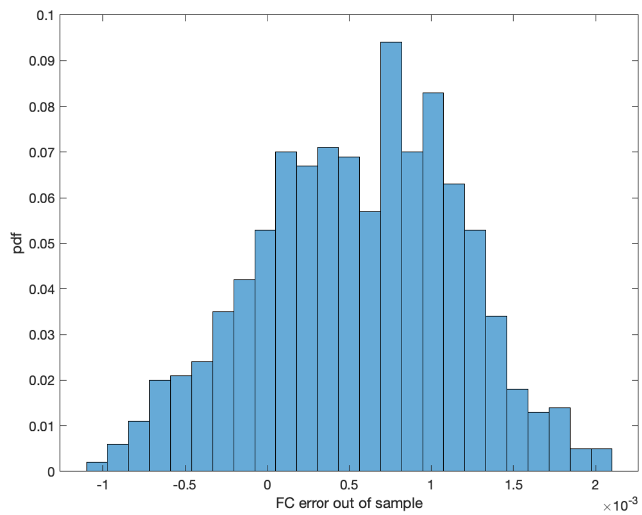

Figure 5: Full Dependent GPCE Method

Compared to the one-period ahead forecast errors the mode of the distribution is shifted to the right with larger errors on average. However, in absolute terms the out-of-sample sample forecast errors are still small showing that the approximate law of motion is accurate over time and forecast errors do not systematically accumulate to large errors. In particular, the average out-of-sample forecast deviates from realized capital by roughly 0.07%. This does not hold for the standard GPCE method. Table 9 documents the average absolute one-period ahead and out-of-sample forecast errors for the Dependent and standard GPCE methods, respectively.

Error\Method	Full Dependent GPCE	Standard GPCE
1-period ahead	$2 \cdot 10^{-4}$	0.0529
out-of-sample	$7 \cdot 10^{-4}$	0.2080

Table 9: Average and maximum forecast errors when evaluating the law of motion out of sample

While average forecast errors between out-of-sample and one-period ahead increase by roughly the same factor the standard GPCE method starts at a significantly higher base as it already produces very large average one-period ahead forecast errors. In particular, forecasts miss realized capital on average by 5.29%. This results in an accumulated average out-of-sample forecast error of 20.8% for the standard GPCE method. To put this into perspective, the percentage standard deviation of the aggregate capital stock in this economy is roughly 0.6%, meaning that the simulated out-of-sample economy moves to aggregate capital stocks significantly and consistently below the interval of realized capital stocks in the proper simulation when using the standard GPCE method. In contrast, the average out-of-sample forecast error of 0.07% of the Dependent GPCE method is still very relative to average variations of the capital stock.

8 GPCE with Time-Varying Base Distributions

Many interesting economic settings are such that the cross-sectional household distribution at times moves far away from the ergodic distribution. Models with financial intermediaries often exhibit strong aggregate non-linearities between the crisis and regular region of the aggregate state space. In that case, instead of using one general base distribution and set of polynomials it is more efficient to choose different bases for different regions in the aggregate state space. In order to assess the accuracy of the dependent GPCE method when projecting on different bases, I solve a model with cyclical earnings risk and different base distributions in recessions and expansions. Evaluating the method's performance when allowing for cyclicality in earnings risk is moreover useful because the objective of this project going forward is to account for the empirically estimated cyclicality of earnings dynamics and revisit the positive conclusions drawn in this paper so far as well as to extend the analysis to normative questions.

The model is identical to the one presented in section 3 apart from the earnings process. The specification of the earnings process changes in that the variance of the persistent component now depends on the aggregate state z_t as in Storesletten et al. (2004).

$$y_t = \exp\left(p_t(z_t) + \epsilon_t\right), \qquad \epsilon_t \underset{iid}{\sim} \mathcal{N}(0, \sigma_\epsilon^2) \tag{8.1}$$

$$p_t = \rho p_{t-1} + \eta_t(z_t), \qquad \eta_t \underset{iid}{\sim} \mathcal{N}(0, \sigma_\eta^2(z_t)) \tag{8.2}$$

Table 10 depicts the calibration of the process which targets the estimates from Storesletten et al. (2004). I use their reported moments for the persistent and transitory components and compute the implied moments of the overall earnings process. I then use GMM to calibrate and discretize the persistent earnings processes in recessions and expansions while ensuring a common earnings grid. The targets are the implied moments for earnings [6]

Parameter	Value	Description
$\sigma_y^2(R)$	0.6	Variance earnings shocks in Recession
$\sigma_y^2(E)$	0.33	Variance earnings shocks in Expansion
σ_ϵ^2	0.22	Variance transitory shocks

Table 10: Earnings process with cyclical variance, targets based on Storesletten et al. (2004)

[6] I use the same discretization method as for the higher-order process but without targeting higher-order moments. See appendix B for more details.

8.1 Evaluating Cyclical Dependent GPCE

I evaluate the solution method again based on forecast errors for aggregate capital. For comparability, I truncate the polynomial at order three as in the main model. Figure 6 shows one-period ahead forecast errors for capital.

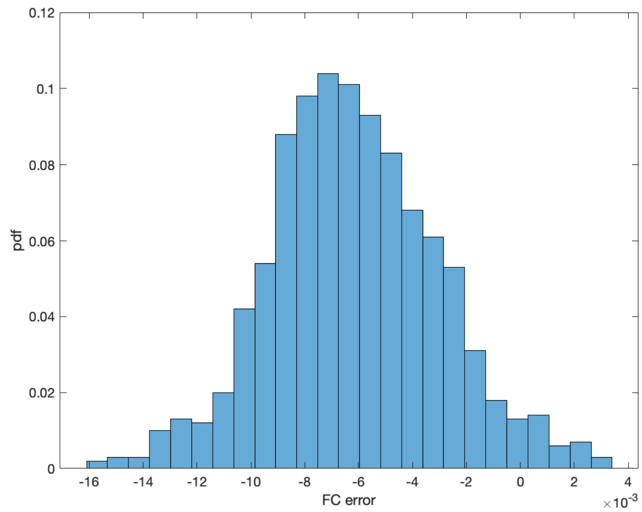

Figure 6: One-period ahead forecasting error distribution with time-varying base distribution

One-period ahead forecast errors are roughly one order of magnitude larger than in the non-cyclical economy in which the base distribution is constant. On average households miss the realized capital stock with their forecasts by roughly 0.5%. While the forecast errors increase relative to the canonical economy this is to be expected as capital is also substantially more volatile with a percentage standard deviation of 3.75%. The intuition is that countercyclical earnings risk amplifies aggregate fluctuations by inducing a countercyclical precautionary savings motive. While I am still working on improving the solution method, this shows that it can be used when the base distribution varies over time.

9 Conclusion

This paper has introduced higher-order earnings dynamics consistent with recent empirical findings into a workhorse heterogeneous-agent real business cycle model. Compared to canonical earnings processes higher-order earnings risk induces larger earnings inequality, particularly through more persistent and lower earnings at the bottom of the distribution. Low-income households at the bottom of the wealth distribution respond by reducing consumption as a share of total resources by roughly 13 percentage points and thereby increase savings. In contrast, wealthier households are not strongly affected by changes in the higher-order moments of earnings shocks as they behave as permanent income consumers. As a result, wealth inequality falls while consumption inequality increases, reinforcing the known issue of generating the degree of wealth dispersion observed in the data. Since effectively only poor households adjust their consumption-savings pattern the aggregate time series dynamics of the model does not change in the presence of higher-order earnings risk, leaving the dynamics of aggregate capital, output and consumption mostly unchanged. However, higher-order earnings risk does imply substantial welfare costs for wealth poor households equivalent to up to 1.7% of permanent consumption. There are two reasons for that. First, average consumption falls for the poor as the precautionary savings motive increases. Second, for poor households resources to spend consist mostly of earnings. Therefore, the higher-order moments of earnings are passed through to consumption dynamics, and particularly excess kurtosis induces larger changes in consumption which lowers welfare.

References

Aiyagari, S. R. (1994) "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, Vol. 109, pp. 659–684.

Auclert, A., B. Bardoczy, M. Rognlie, and L. Straub (2021) "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*.

Bach, L., L. E. Calvet, and P. Sodini (2020) "Rich Pickings? Risk, Return, and Skill in Household Wealth," *American Economic Review*, Vol. 110, pp. 2703–2747.

Bewley, T. (1977) "The permanent income hypothesis: A theoretical formulation," *Journal of Economic Theory*, Vol. 16, pp. 252–292.

Blanchard, O. J. (1985) "Debt, Deficits, and Finite Horizons," *Journal of Political Economy*, Vol. 93, pp. 223–247.

Boppart, T., P. Krusell, and K. Mitman (2018) "Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative," *Journal of Economic Dynamics and Control*, Vol. 89, pp. 68–92.

Brunnermeier, B. M. K. and Y. Sannikov (2014) "A Macroeconomic Model with a Financial Sector," *American Economic Review*, Vol. 104, pp. 379–421.

Busch, C. and A. Ludwig (2020) "Higher-Order Income Risk over the Business Cycle."

Cagetti, M. and M. De Nardi (2006) "Entrepreneurship , Frictions , and Wealth Mariacristina De Nardi," *Journal of Political Economy*, Vol. 114, pp. 835–870.

Carroll, C. D., J. Slacalek, and K. Tokuoka (2015) "Buffer-stock saving in a Krusell-Smith world," *Economics Letters*, Vol. 132, pp. 97–100.

Carroll, C., J. Slacalek, K. Tokuoka, and M. N. White (2017) "The distribution of wealth and the marginal propensity to consume," *Quantitative Economics*, Vol. 8, pp. 977–1020.

Chang, Y., S. B. Kim, K. Kwon, and R. Rogerson (2019) "2018 Klein Lecture: Individual and Aggregate Labor Supply in Heterogeneous Agent Economies With Intensive and Extensive Margins," *International Economic Review*, Vol. 60, pp. 3–24.

Civale, S., L. Diez-Catalan, and F. Fazilet (2017) "Discretizing a Process with Non-zero Skewness and High Kurtosis."

De Nardi, M., G. Fella, and G. Paz-pardo (2018) "Nonlinear Household Earnings Dynamics, Self-insurance, and Welfare," *NBER Working Paper*.

Den Haan, W. J. (1996) "Heterogeneity, Aggregate Uncertainty, and the Short-Term Interest Rate," *Journal of Business & Economic Statistics*, Vol. 14, pp. 399–411.

——— (1997) "Solving dynamic models with aggregate shocks and heterogeneous agents," *Macroeconomic Dynamics*, Vol. 1, pp. 355–386.

——— (2010) "Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents," *Journal of Economic Dynamics and Control*, Vol. 34, pp. 79–99.

Fagereng, A., M. B. Holm, and G. J. J. Natvik (2020) "MPC Heterogeneity and Household Balance Sheets," *SSRN Electronic Journal*.

Fernández-Villaverde, J., S. Hurtado, and G. Nano (2020) "Financial Frictions and the Wealth Distribution," *NBER Working Paper*.

Gautschi, W. (1982) "On Generating Orthogonal Polynomials," *SIAM Journal on Scientific and Statistical Computing*, Vol. 3, pp. 289–317.

Golosov, M., M. Troshkin, and A. Tsyvinski (2016) "Redistribution and social insurance," *American Economic Review*, Vol. 106, pp. 359–386.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2015) "What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk?," *SSRN Electronic Journal*.

Guvenen, F., S. Ozkan, and J. Song (2014) "The nature of countercyclical income risk," *Journal of Political Economy*, Vol. 122, pp. 621–660.

Hubmer, J., P. Krusell, and A. A. Smith (2021) "Sources of us wealth inequality: Past, present, and future," *NBER Macroeconomics Annual*, Vol. 35, pp. 392–455.

Huggett, M. (1993) "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of Economic Dynamics and Control*, Vol. 17, pp. 953–969.

Krueger, D. and A. Ludwig (2016) "On the Optimal Provision of Social Insurance: Progressive Taxation versus Education Subsidies in General Equilibrium," *forthcoming: Journal of Monetary Economics*.

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

Krusell, P. and A. A. Smith (1998) "Income and wealth heterogeneity in the macroeconomy," *Journal of Political Economy*, Vol. 106, pp. 867–896.

Maliar, L., S. Maliar, and F. Valli (2010) "Solving the incomplete markets model with aggregate uncertainty using the Krusell-Smith algorithm," *Journal of Economic Dynamics and Control*, Vol. 34, pp. 42–49.

Proehl, E. (2017) "Approximating Equilibria with ex-post Heterogeneity and Aggregate Risk."

Quadrini, V. (2000) "Entrepreneurship, Saving, and Social Mobility," *Review of Economic Dynamics*, Vol. 3, pp. 1–40.

Reiter, M. (2010) "Solving the incomplete markets model with aggregate uncertainty by backward induction," *Journal of Economic Dynamics and Control*, Vol. 34, pp. 28–35.

Schaab, A. (2021) "Micro and Macro Uncertainty," *Job Market Paper*.

Storesletten, K., C. I. Telmer, and A. Yaron (2004) "Cyclical Dynamics in Idiosyncratic Labor Market Risk ABI/INFORM Global pg. 695," *The Journal of Political Economy*, Vol. 112.

Winberry, T. (2018) "A method for solving and estimating heterogeneous agent macro models," *Quantitative Economics*, Vol. 9, pp. 1123–1151.

Zheng, M., X. Wan, and G. E. Karniadakis (2015) "Adaptive multi-element polynomial chaos with discrete measure: Algorithms and application to SPDEs," *Applied Numerical Mathematics*, Vol. 90, pp. 91–110.

mrohorolu, A. (1989) "Cost of Business Cycles with Indivisibilities and Liquidity Constraints," *Journal of Political Economy*, Vol. 97, pp. 1364–1383.

A Generating the Polynomials

Given a base random variable ν the corresponding polynomials are generated with the three-term recurrence relation given by

$$\Psi_{i+1}(\nu) = (\nu - \gamma_i)\Psi_i(\nu) - \beta_i \Psi_{i-1}(\nu), \qquad i \geq 1$$

where $\Psi_0(\nu) = 0$. The sequence is initialized with $\Psi_1(\nu) = 1$. As discussed the first polynomial is always constant and equal to 1 which yields that the first coefficient $\alpha_{t,1}$ is equal to the mean of the distribution. The parameters $\gamma_i, \beta_i \in \mathbb{R}$ can be generated in different ways as discussed in Zheng et al. (2015). Following Proehl (2017) I use the Stieltjes method explained in detail in Gautschi (1982).

$$\gamma_i = \frac{\langle \Psi_i, \nu\Psi_i \rangle}{\langle \Psi_i, \Psi_i \rangle}$$
$$\beta_i = \frac{\langle \Psi_i, \Psi_i \rangle}{\langle \Psi_{i-1}, \Psi_{i-1} \rangle}$$

where $\langle :, : \rangle$ is the standard inner product in L^2 with respect to the base distribution ν.

B Discretizing the Higher-Order Process

B.1 Targets and Calibrated Parameters

The earnings process is characterized as follows

$$y_t = \exp\left(p_t + \epsilon_t\right), \qquad \epsilon_t \underset{iid}{\sim} \mathcal{N}(0, \sigma_\epsilon^2)$$
$$p_t = \rho p_{t-1} + \eta_t, \qquad \eta_t \underset{iid}{\sim} \begin{cases} \mathcal{N}(\mu_1, \sigma_1^2) \text{ with } p_1 \\ \mathcal{N}(\mu_2, \sigma_2^2) \text{ with } 1 - p_1 \end{cases}$$

I use GMM to solve for the calibration of the above process as well as the grid for the persistent component p_t to match a set of moments of the earnings process. Let $\vec{p}$ denote the grid of persistent earnings then the set of parameters to choose is $\theta = (p_1, \mu_1, \mu_2, \sigma_1, \sigma_2, \sigma_\epsilon, \vec{p})$. Given the standard deviation σ_ϵ the transitory process is discretized with Tauchen [7]. I first use a global solver and afterwards improve on the initial solution with a more powerful

[7] I discretize the transitory component into 3 possible values and the persistent component into a 7 point Markov chain. This generates a 21 point Markov chain for overall earnings

local solver. The local solver targets the same set of moments but only chooses the grid for persistent earnings $\theta_{local} = \vec{p}$, not the calibration parameters for the process.
Let x denote log-earnings and $\Delta_j x = x_t - x_{t-j}$ the j-year growth rate of log-earnings. Table 11 shows the set of targets as well as corresponding moments of the discretized earnings process.

	Targets	Discretized process
Variance(x)	0.58	0.59
Variance($\Delta_1 x$)	0.23	0.19
Skewness($\Delta_1 x$)	-1.35	-1.39
Kurtosis($\Delta_1 x$)	17.8	15.72
Variance($\Delta_5 x$)	0.46	0.51
Skewness($\Delta_5 x$)	-1.01	-1.25
Kurtosis($\Delta_5 x$)	11.55	9.88

Table 11: Unconditional moments of log earnings for higher-order earnings process, targets from Guvenen et al. (2015)

B.2 Moments of the Discretized Process

This section briefly derives the targeted moments of the discretized earnings process given a set of parameters θ. Let $\mathcal{F}$ and $\mathbf{f}$ denote the cdf and pdf of the Gaussian mixture distribution of the innovations of the persistent component, given a set of calibration parameters. Further, let $\vec{p}$ denote the grid of persistent earnings and n_p, n_ϵ the grid sizes of persistent and transitory components, respectively.
Based on $\vec{p}$ one generates another grid $\vec{d}$ across the real numbers with buckets for each persistent earnings state given by

$$\vec{d}(i) = \frac{\vec{p}(i-1) + \vec{p}(i)}{2}, \qquad i \geq 2$$

where $\vec{d}(1) = -\infty$ and $\vec{d}(n_p) = \infty$. Then the transition matrix T_p for persistent earnings can be generated in the following way

$$T_p(r,c) = \mathcal{F}(\vec{d}(c+1) - \rho\vec{d}(r)) - \mathcal{F}(\vec{d}(c) - \rho\vec{d}(r)).$$

Given the transition matrix for persistent earnings, it is straight forward to derive the corresponding invariant distribution of persistent earnings π_p. The transitory component is discretized with Tauchen's method which yields the corresponding grid $\vec{\epsilon}$, transition matrix

T_ϵ, and invariant distribution π_ϵ. The possible values and, thus, the grid for overall earnings $\vec{y}$ follows from all possible combinations of persistent and transitory component according to the specified process $y(p, \epsilon) = \exp(p + \epsilon)$. The transition probabilities for earnings T_y is given by

$$T_y(y'(p', \epsilon')|y(p, \epsilon)) = T_p(p'|p) \cdot T_\epsilon(\epsilon'|\epsilon)$$

As with the persistent component the invariant distribution of earnings π_y can be derived by iterating over the transition matrix T_y until convergence. The grid $\vec{y}$, transition matrix T_y and invariant distribution π_y fully characterize the AR(1) process for earnings and are sufficient to derive the resulting moments.