Deep Learning in the Sequence Space*

Marlon Azinovic-Yang†
University of North Carolina at Chapel Hill

Jan Žemlička‡
University of Zurich and Swiss Finance Institute

First version: September 9, 2025
This version: March 13, 2026

Abstract

We develop a deep learning algorithm for constructing globally accurate approximations to functional rational expectations equilibria of dynamic stochastic economies in the sequence space. We use deep neural networks to parameterize key equilibrium objects, such as policies or prices, as functions of truncated histories of exogenous shocks. We train the neural networks to satisfy equilibrium conditions along simulated paths of the economy. We illustrate the performance of our method in three environments: (i) a high-dimensional overlapping generations economy with multiple sources of aggregate risk; (ii) an economy with heterogeneous households and firms facing uninsurable idiosyncratic risk and large shocks to idiosyncratic and aggregate volatility; and (iii) a stochastic life-cycle economy with a continuous asset choice and a discrete early-retirement choice that induces local convexities in the continuation values of working-age cohorts. We also propose practical neural policy architectures that guarantee monotonicity of predicted policies, enabling the endogenous grid method to simplify parts of the algorithm. We achieve high precision throughout, with the mean error in equilibrium conditions below 0.2%.

JEL classification: C61, C63, C68, D52, E32.

Keywords: deep learning, deep neural networks, sequence space, heterogeneous firms, heterogeneous households, overlapping generations, life-cycle, global solution method, discrete choice.

1 Introduction

We bring deep learning solution methods to the sequence space and propose a new global method for computing equilibria in economies with aggregate risk.[1] Exploiting the ergodicity property of a large

*This paper has greatly benefited from numerous discussions with Felix Kubler, Jonathan Payne, Simon Scheidegger and Yucheng Yang. We are also grateful for the comments of participants at seminars at the University of Zurich, UNC Chapel Hill, Princeton University, City University of New York, and the 2025 Torino Conference on Machine Learning for Economics and Finance as well as for helpful discussions with Jaroslav Borovička, Anusha Chari, Jaden Chen, Jeppe Druedahl, Lutz Hendricks, William Jungerman, Yasutaka Koike-Mori, Lilia Maliar, Stanislav Rabinovich, Jacob Røpke, and Can Tian. Žemlička gratefully acknowledges support from the Swiss National Science Foundation and the kind hospitality of the Department of Economics at Princeton University.

†Email: marlonay@unc.edu

‡Email: jan.zemlicka@df.uzh.ch

[1] We use the expression "sequence space" liberally, referring to the space of truncated sequences of aggregate shocks, not (only) the space of perfect foresight sequences. To be more precise, one could alternatively refer to our approach as a moving-average rather than a sequence-space approach.

class of dynamic economies, we approximate the aggregate state vector using a truncated sequence of aggregate shocks. We use deep neural networks to parameterize the mapping from truncated histories of aggregate shocks to equilibrium objects of interest, and optimize them using stochastic gradient descent to minimize errors in the equilibrium conditions over the simulated ergodic set of the economy.

We make four distinct contributions to the literature. First, we show how deep learning can be used to compute global approximations to equilibria of dynamic stochastic economies in the sequence space. Second, we show how to construct neural network architectures that predict approximate policy functions, which are guaranteed to satisfy ex-ante known monotonicity and concavity properties. Third, based on our monotonicity-preserving architectures, we show how to obtain supervised learning targets for policy and price functions, using the endogenous gridpoint method of Carroll (2006). Fourth, we illustrate the rich applicability and accuracy of our method by solving three models, each featuring a different computational challenge. We first solve an overlapping generations model with portfolio choice and regime switching. We then proceed to compute a global solution to an economy with a continuum of heterogeneous firms and a continuum of heterogeneous households. To our knowledge, we are the first to do so in an economy featuring incomplete markets and household heterogeneity in the spirit of Imrohoroğlu (1989)-Bewley (1977)-Huggett (1993)-Aiyagari (1994)-Krusell and Smith (1998), together with firm heterogeneity in the spirit of Bloom et al. (2018) and Khan and Thomas (2008). Lastly, we apply our method to solve an overlapping generations model with intra- and intergenerational risk and a discrete early-retirement choice. The discrete choices lead to local convexities in the continuation values of working-age cohorts, rendering their Euler equations insufficient for uniquely characterizing their optimal savings policies.

To describe the main idea of our method, we now take the canonical Krusell and Smith (1998) model as an example. Let z_t denote the realization of stochastic TFP in period t, and let $z^t := (z_0, z_1, \ldots, z_t)$ denote the history of realized shocks. In an infinite horizon model the history of shocks becomes infinitely long. The state of the economy is the current level of TFP z_t, together with the cross-sectional distribution of households over wealth and productivity. Let e_t^i denote the idiosyncratic productivity of the household i in period t, let k_t^i denote their capital holdings, and let $\mu_t = \mu_t(e, k)$ denote the cross-sectional distribution. Since the distribution, even when discretized into a histogram as in Young (2010), is a high-dimensional object, the presence of the distribution in the state vector poses a long-standing challenge for solution methods.[2] The approach of Krusell and Smith (1998) is to approximate the cross-sectional distribution μ_t with a low-dimensional statistic such as aggregate capital K_t. Instead of approximating the state vector using its moments or a histogram, we build on the seminal perturbation approaches in the sequence space developed in Boppart et al. (2018) and Auclert et al. (2021) and approximate the aggregate state vector using *sequence-space* objects. In particular, we propose to approximate the aggregate state of the economy by a truncated history of exogenous aggregate shocks. In the context of the Krusell and Smith (1998) model, for example, let $z_T^t = (z_{t-T}, z_{t-T+1}, \ldots, z_{t-1}, z_t)$ denote the last T realizations of TFP. Instead of solving for equilibrium functions $f^{\text{approx}}(z_t, K_t) \approx f(z_t, \mu_t)$, we propose to solve for equilibrium functions $f^{\text{approx}}(z_T^t) \approx f(z_t, \mu_t)$. For any $T < \infty$ our approach implies a truncation

[2] See, for example, Moll (2025).

error. However, in economies featuring ergodic dynamics, this truncation error converges to zero as $T \to \infty$. Hence, we can approximate the information contained in the aggregate state vector arbitrarily well by choosing T sufficiently large. Consequently, our function approximators need to be able to handle large T, implying a high-dimensional input z_t^T, potentially even higher dimensional than a histogram of the wealth distribution. We will refer to our approach, which uses a truncated history of aggregate shocks as input as *sequence-space approach* and to the standard approach (using, for example, the wealth distribution or a lower dimensional sufficient statistic such as standard moments, as input to the equilibrium functions) as *state-space approach*.

We use neural networks to approximate all equilibrium functions that need to be solved for. Like polynomials, neural networks are universal function approximators (Hornik et al., 1989). More importantly however, neural networks have shown great promise in ameliorating the curse of dimensionality (Bellman, 1961) associated with high-dimensional function-approximation problems, such as solving partial differential equations (see, *e.g.* Grohs et al., 2018; Jentzen et al., 2018) or computing equilibria in dynamic economies (see, *e.g.* Azinovic et al., 2022; Maliar et al., 2021; Kahou et al., 2021; Gu et al., 2023).[3] The sequence-space formulation has two main advantages in the context of deep learning solution methods. First, the sequence-space formulation is potentially more parsimonious, especially in economies with rich cross-sectional heterogeneity.[4] Second, using truncated shock histories as the only network input mitigates a dangerous feedback loop present in simulation-based deep learning solution methods. While the distribution of shock histories is exogenous and fixed,[5] the distribution of endogenous states moves during the training as the algorithm updates the approximate policy function. As discussed in Azinovic et al. (2022), a large update to the policy function might shift the distribution of training states into areas that have been previously only sparsely covered, causing large sample error, and correspondingly large gradients, potentially inducing further jumps in the policy function. While our algorithm still relies on simulating endogenous states as an input for evaluating loss function[6] those do not enter as an input into neural networks parameterizing the core equilibrium objects of the economy.

The second contribution of our paper is that we show how to construct neural network architectures that guarantee monotonicity or even concavity of the resulting approximator with respect to a chosen set of inputs. Rather than directly approximating individual policy function using a neural network that takes the aggregate and idiosyncratic state as input,[7] we build on the operator learning approach of Zhong (2023). An operator network maps the sequence of aggregate shocks to policy *functions*. The predicted policy functions then map idiosyncratic state variables into the individual policy choices. Therefore, instead of having to repeatedly evaluate the neural network for many combinations of aggregate and idiosyncratic states, the neural network predicts a single policy function for each aggregate state. The predicted policy function is then evaluated for all

[3] Section 2 provides a more detailed discussion of the related literature in economics.

[4] e.g. The state-space formulation of the workhorse HANK economy features a state vector that includes a cross-sectional distribution of households over idiosyncratic shocks as well as liquid and illiquid assets. A tensor-product discretization of such a distribution might easily generate tens of thousands of aggregate state variables, whereas a highly accurate sequence-space representation could be obtained using a few hundred of shock lags at quarterly frequency.

[5] Because the distribution of fundamental shocks is just a primitive of the economic model at hand.

[6] e.g. for evaluation of marginal product of capital

[7] As in Azinovic et al. (2022) or Maliar et al. (2021).

idiosyncratic states of interest. For the Krusell and Smith (1998) economy, for example, the neural network predicts a consumption *function* based on the realized history of aggregate shocks, which then takes only individual productivity and asset holdings as input. Additionally, ensuring monotonicity allows us to implement parts of the algorithm using the endogenous grid method (EGM) (Carroll, 2006). The use of EGM is particularly advantageous when tackling models where agents face a mix of discrete and continuous choices since it enables the use of an upper envelope algorithm (Druedahl and Jørgensen, 2017; Dobrescu and Shanker, 2023).

2 Literature

The idea of approximating the endogenous aggregate state by a truncated sequence of aggregate shocks was first introduced by Chien et al. (2011) in the context of an endowment economy with heterogeneous agents and with a single aggregate shock that takes two discrete values. In their setting, they find that considering the history of the past five aggregate shocks is enough to yield a high-accuracy approximation. In the context of production economies with aggregate risk Lee (2025) introduces the repeated transition method (RTM). The underlying idea of the RTM is that, in an ergodic economy, all possible states will be realized along a sufficiently long simulation path. Lee (2025) uses this fact to construct the continuation value function for the Bellman equations of the agents in the economy. The RTM iterates over simulating a long path of the economy. In each iteration it constructs the continuation values by using the previously simulated path as a look-up table. Our contribution relative to this literature is generality. While those two papers also develop sequence-based global solution methods, the applicability of these methods has limitations that do not pose a significant challenge for our method. Namely, the algorithm of Chien et al. (2011) struggles to resolve economies featuring stronger persistence, such as production economies, because strong persistence implies the need to keep track of a long history of shocks. The RTM approach of Lee (2025) can handle a rich set of environments, including production economies; however, its main bottleneck lies in handling economies with a large number of different aggregate shocks. At each iteration the RTM algorithm has to construct the continuation value function for every *period*. To do so, the algorithm has to split the simulated path into partitions sorted by the realized value of aggregate shock, and then to search through each partition for a realization where the aggregate endogenous state (e.g. wealth distribution) is closest to the *next period* distribution with respect to some norm. This structure implies that the RTM algorithm works best in environments with a small number of discrete aggregate shocks. Complex stochastic processes that cannot be approximated using a parsimonious Markov chain pose a challenge for RTM, as the length of simulated path required for ergodic representation grows with the cardinality of shock discretization. In contrast, we demonstrate that our method can efficiently solve rich economies featuring production and complex driving stochastic processes. Besides that, our algorithm can handle discrete as well as continuous shocks. Furthermore, neural networks can learn from a large number of simultaneous trajectories, allowing for significant acceleration on modern massively parallel computing architectures.[8]

[8]Global solution methods to solve models on the simulated paths using state-space approaches include Brumm and Hußmann (2025), Levintal (2018), and Azinovic et al. (2023)

Instead of using the truncated history of exogenous shocks as a summary statistic for the aggregate state, recent work by Yang et al. (2025) and Guarda (2025) uses short histories of endogenous aggregate variables in the spirit of Krusell and Smith (1998). Yokomoto (2025) implements the Krusell and Smith (1998) algorithm using neural networks as function approximators. Yang et al. (2025) and Yokomoto (2025) use prices as their summary statistic for the endogenous aggregate state, allowing for fast computation of market clearing prices in the simulation step. Yang et al. (2025) update policies with a policy gradient algorithm to maximize simulated lifetime rewards, while Yokomoto (2025) updates policies using the Bellman equation. Households' decisions depend on the endogenous aggregate state only through prices. Hence, households directly track the aggregate variable that enters into their optimization problem. In Yang et al. (2025) very short histories are sufficient to forecast conditional future prices well, keeping the state-space for households low-dimensional and thereby addressing the challenges raised in Moll (2025).[9] Although we focus on histories of exogenous shocks, our method can easily be extended to additionally use histories of endogenous aggregate variables, such as prices. In contrast to Yang et al. (2025) our approach works directly with equilibrium conditions and hence allows us to obtain supervised learning targets for household policies. Furthermore, our approach applies to models where the individual problem depends directly on a high dimensional aggregate state and not only through a few low-dimensional aggregate statistics, such as prices.[10]

Perturbation methods in the sequence space have become the method of choice for solving heterogeneous agent models following the game-changing work by Boppart et al. (2018) and Auclert et al. (2021). Boppart et al. (2018) and Auclert et al. (2021) introduce local methods, which linearize heterogeneous agent models with respect to aggregates, in the sequence space.[11] These methods, together with the accompanying libraries,[12] have substantially increased the tractability of heterogeneous agent models with aggregate risk, rendering a wide set of previously intractable models computationally tractable. Our method complements this literature by providing a global method, which is hence applicable for models with larger shocks and without certainty equivalence with respect to aggregate risk. The cost of having a global method is that the runtime is substantially longer in comparison to local methods, such as the sequence-space Jacobian method by Auclert et al. (2021).

The toolbox that enables us to approximate functions on the very high dimensional domain of long histories of economic shocks is deep learning. Deep learning as a tool to compute equilibria in economic models dates back to Duffy and McNelis (2001), who use a neural network to solve a stochastic growth model using a parameterized expectations algorithm. Norets (2012) makes use of neural networks in the context of discrete-state dynamic programming. Closer to this paper Azinovic et al. (2022) and Maliar et al. (2021) use neural networks to approximate equilibrium price and policy functions, and train them to satisfy equilibrium conditions, such as first-order optimality conditions, Bellman equations, and market clearing conditions along the simulated paths

[9] Applications in Yang et al. (2025) include the Krusell and Smith (1998) model, a Huggett (1993) model with aggregate risk, and a one asset HANK model along the lines of Auclert et al. (2021).

[10] See, for example, Payne et al. (2024).

[11] See also Reiter (2009) and Bayer and Luetticke (2020) for linearization in aggregate variables in the state space.

[12] See `https://github.com/shade-econ/sequence-jacobian` for a user friendly library for the method in Auclert et al. (2021).

of the economy. Azinovic and Zemlicka (2024) extend these methods by introducing market clearing layers[13] and step-wise model transformations to robustly solve models with portfolio choice. Kase et al. (2023) show how to use deep learning to compute the solutions of economic models for ranges of economic parameters, speeding up the maximum likelihood estimation of these parameters. Han et al. (2022) and Kahou et al. (2021) introduce symmetry preserving neural network architectures and Valaitis and Villa (2024) use deep learning within a parameterized expectations algorithm. Fernández-Villaverde et al. (2023) use neural networks for a nonlinear forecasting rule within a Krusell and Smith (1998) algorithm. Kahou et al. (2022) investigate the relationship between transversality conditions and solutions found by a deep-learning based algorithm. Druedahl and Ropke (2026) use deep learning to compute optimal choices in finite-horizon life-cycle models with up to eight durable goods. Druedahl et al. (2026a) show how the methodology of Druedahl and Ropke (2026) can be applied to solve non-convex life-cycle models featuring a mixture of discrete and continuous choices. We view our discrete-choice algorithm introduced in section 6 as complementary to Druedahl et al. (2026a): we focus on general equilibrium setups while keeping the optimization problem of an individual agent comparatively simple, whereas Druedahl et al. (2026a) solve complex individual optimization problems with many states and choices in partial equilibrium. In recent work Druedahl et al. (2026b) build on our approach to construct a global solution to a canonical HANK economy with aggregate risk using deep learning in the sequence space. They study the effect of aggregate uncertainty on fiscal multipliers and impulse responses to aggregate shocks.

Zhong (2023) introduces operator learning, which we build on in section 5. While previous deep learning approaches used neural networks to directly parameterize the mapping from a combination of an aggregate and idiosyncratic state to individual choices, Zhong (2023) uses neural networks to approximate the mapping from an aggregate state into a function which then returns individual choices as a function of individual state. We extend the operator learning method of Zhong (2023) by showing how to encode monotonicity or concavity of the predicted functions into the network architecture. Sun (2025a) uses neural networks to approximate continuation values in combination with otherwise standard methods. Neural network based solution methods in continuous time are developed or applied in Duarte (2018), Sauzet (2021), Gopalakrishna (2021), Gu et al. (2023), Gopalakrishna et al. (2024) and Payne et al. (2024).[14]

Our contribution relative to the existing literature on deep learning solution methods is twofold. First, we show that deep learning can be efficiently used to construct global approximations to the functional rational expectations equilibria in the sequence space. Second, we construct shape-preserving neural network operator network architectures that allow us to incorporate known shape properties of certain equilibrium functions (e.g. monotonicity and concavity of consumption function in canonical consumption-savings problems) directly into the neural network structure.[15] The ability

[13] Market clearing layers are neural network architectures, designed such that a prediction by the neural network is always consistent with market clearing.

[14] Additional applications of deep learning based solution methods in variety of settings include Folini et al. (2025) and Friedl et al. (2023) (climate economics), Carvalho et al. (2025) (production networks), Kase et al. (2025) (HANK with financial frictions), Duarte et al. (2021) (partial equilibrium finite horizon model with rich asset choice), Bretscher et al. (2022) (international business cycle), Sun (2025b) (spatial economics), Jungerman (2023) (monopsony power in the labor market), Adenbaum et al. (2024) (Bewley (1977) type economy), and Reiter (2025) (hybrid solution approach for OLG).

[15] The idea of shape-preservation in functional approximation methods for dynamic economies was introduced by

to guarantee monotonicity of the consumption function allows us to apply the method of endogenous gridpoints of Carroll (2006) and hence simplify the solution algorithm for a large class of problems featuring optimization problems that are amenable to EGM.

3 Algorithm and illustrative application

3.1 Model

We first explain our proposed algorithm by applying it to the simplest workhorse economy with aggregate risk, the stochastic growth model of Brock and Mirman (1972). This model can be accurately solved using conventional methods, so this section does not aim to claim a computational achievement of our method. Instead, it serves to lay out the proposed algorithm as transparently as possible. In sections 4, 5, and 6, we then apply our algorithm in settings that pose a substantial challenge for existing numerical methods.

Time is infinite and discrete, $t = 0, 1, \ldots$. We consider an infinitely lived representative household that derives utility $u(C_t)$ from consumption C_t. The representative household owns the capital stock K_t and inelastically supplies a constant amount of efficient units of labor $L = 1$ at the equilibrium wage w_t. The single good in the economy is produced by a representative firm endowed with a Cobb-Douglas production technology

$$Y_t = A_t K_t^{\alpha} L^{1-\alpha}, \tag{1}$$

where A_t denotes stochastic total factor productivity, that evolves according to an AR(1) process

$$\log(A_t) = \rho^A \log(A_{t-1}) + \sigma^A \epsilon_t^A, \tag{2}$$

where $\epsilon_t^A \sim \mathcal{N}(0, 1)$. The firm rents capital K_t and efficient units of labor L on competitive spot markets and pays a rental rate $r_t^K = \alpha A_t K_t^{\alpha-1} L^{\alpha}$ on capital and a wage $w_t = (1-\alpha) A_t K_t^{\alpha} L^{-\alpha}$ for efficient units of labor. In each period, the representative household chooses how much to consume and how much to invest in capital. Capital evolves according to $K_{t+1} = (1-\delta)K_t + I_t$, where I_t denotes the investment in capital and δ denotes its depreciation rate. The households' problem is given by

$$\max_{\{K_t\}_{t=1}^{\infty}} \mathrm{E}\left[\sum_{t=0}^{\infty} \beta^t u(C_t)\right] \tag{3}$$

subject to :

$$C_t = L w_t + (1 - \delta + r_t^K) K_t - K_{t+1}.$$

The Euler equation, which characterizes the equilibrium together with a transversality condition, is

Judd and Solnick (1994), and further developed by Cai and Judd (2012).

given by

$$u'(C_t) = \mathrm{E}\left[\beta u'(C_{t+1})(1-\delta+r^K_{t+1})\right] \tag{4}$$

$$\Leftrightarrow 0 = \frac{(u')^{-1}\left(\mathrm{E}\left[\beta u'(C_{t+1})(1-\delta+r^K_{t+1})\right]\right)}{C_t} - 1. \tag{5}$$

Where the Euler equation (5) is rearranged, such that errors in the optimality condition can be interpreted as relative consumption errors. This allows for an implicit, yet interpretable accuracy measure (see Judd, 1998). Traditionally, the model would be solved using recursive methods, where the state of the economy $\mathbf{X}^{\text{state}}_t := [A_t, K_t] \in \mathbb{R}^2$ is given by the current level of TFP together with the current capital stock. The solution to the model is a policy function $\pi(\mathbf{X}^{\text{state}}_t) = K_{t+1}$, which maps the state of the economy to the endogenous variables of interest, in this case the capital stock for the next period.[16]

3.2 Algorithm

3.2.1 Main idea

Our algorithm uses a deep neural network to approximate the policy function. Let $\mathcal{N}_{\boldsymbol{\rho}}$ denote a neural network with trainable parameters $\boldsymbol{\rho}$. Following the *Deep Equilibrium Nets* algorithm (or similar algorithms often used in the literature) the goal is to approximate the policy function by a neural network. This would mean to find neural network parameters $\boldsymbol{\rho}$, such that for all states

$$\mathcal{N}_{\boldsymbol{\rho}}(\mathbf{X}^{\text{state}}_t) \approx \pi(\mathbf{X}^{\text{state}}_t) = K_{t+1}. \tag{6}$$

The main idea of the algorithm we propose in this paper is to train the neural network to predict policies not based on $\mathbf{X}^{\text{state}} = [A_t, K_t] \in \mathbb{R}^2$ but instead purely based on a truncated history of shocks, $\mathbf{X}^{\text{seq},T} := [A_{t-T+1}, A_{t-T+2}, \ldots, A_{t-1}, A_t] \in \mathbb{R}^T$, such that

$$\mathcal{N}_{\boldsymbol{\rho}}(\mathbf{X}^{\text{seq},T}_t) \approx \pi(\mathbf{X}^{\text{state}}_t) = K_{t+1}. \tag{7}$$

At first sight, using the history of shocks as an input may look like a step in the wrong direction for two reasons: first, for $T < \infty$ the sequence of shocks is only an approximately sufficient statistic for the state of the economy. Hence, by truncating the history, we introduce a limit to the precision that even a perfectly trained neural network could attain. Second, the dimensionality of $\mathbf{X}^{\text{seq},T}$ is sometimes larger than the dimensionality of $\mathbf{X}^{\text{state}}$.[17] Therefore, the sequence space approach would not work well together with approximation methods, for which a low-to-medium dimensional domain is pivotal, such as, for example, (adaptive) sparse grids.[18] As we illustrate in this paper, however, the sequence approximation to the endogenous aggregate state is promising when used together with deep neural networks as function approximators.

[16] There are other equivalent policies, like a policy for investment or consumption, which together the budget constraint and the law of motion for capital lead to the same choice for capital stock in the next period.

[17] For example in this simple stochastic growth economy.

[18] See Krueger and Kubler (2004) for sparse grids and Brumm and Scheidegger (2017) for adaptive sparse grids.

3.2.2 Sufficiency of the truncated history of shocks

For our method to work, it must be that the truncated sequence of shocks is an approximate sufficient statistic for the endogenous aggregate state of the economy, in this case capital K_t. This can only be the case if the influence of capital at period $t-T$, K_{t-T}, on capital in period t, K_t, is vanishing as $T \to \infty$, *i.e.* our method relies on the ergodicity property of the underlying economy.

To see that this condition indeed holds in our stochastic growth economy, we now consider a special case with full depreciation $\delta = 1$ and logarithmic utility $u(C) := \log(C)$. In this case, the model admits a closed form solution with $K_{t+1} = \alpha\beta A_t K_t^\alpha$. Taking log on both sides, we obtain $\log(K_{t+1}) = \log(\alpha\beta) + \log(A_t) + \alpha \log(K_t)$. Iterating forward, we obtain

$$\begin{aligned}
\log(K_t) &= \log(\alpha\beta) + \log(A_{t-1}) + \alpha \log(K_{t-1}) \\
&= \log(\alpha\beta) + \log(A_{t-1}) + \alpha(\log(\alpha\beta) + \log(A_{t-2}) + \alpha \log(K_{t-2})) \\
&= \log(\alpha\beta) + \log(A_{t-1}) + \alpha(\log(\alpha\beta) + \log(A_{t-2}) + \alpha(\log(\alpha\beta) + \log(A_{t-3}) + \alpha \log(K_{t-3}))) \\
&= \cdots = \underbrace{f(\alpha, \beta, A_{t-1}, A_{t-2}, A_{t-3}, \ldots, A_{t-T})}_{\text{function of } \alpha,\ \beta \text{ and the last } T \text{ productivity values}} + \underbrace{\alpha^T \log(K_{t-T})}_{\text{truncation error}}, \qquad (8)
\end{aligned}$$

where the function f depends only on the parameters α, β and the truncated history of the last T productivity values. The value of capital at $t-T$ only affects the value of capital with a coefficient of α^T, where $\alpha < 1$ is the share of capital in production, typically around 0.3. The truncation error therefore vanishes exponentially at the rate α.

As an alternative to using lagged productivity levels, we could use a truncated history of innovations as the input to the neural network.[19] Iterating the AR(1) implies

$$\begin{aligned}
\log(A_t) &= \rho^T \log(A_{t-T}) + \sigma \sum_{j=0}^{T-1} \rho^j \epsilon_{t-j} \qquad (9) \\
&=: g(\rho, \sigma, \epsilon_{t-T+1}, \ldots, \epsilon_t) + \underbrace{\rho^T \log(A_{t-T})}_{\text{truncation error}}. \qquad (10)
\end{aligned}$$

Hence, for large enough T and $|\rho| < 1$ the truncated history of innovations is an approximate sufficient statistic for the history of productivity values, which in turn are an approximate sufficient statistic for the endogenous aggregate state.

The same argument can also be made in a single step by observing that

$$\begin{bmatrix} \log(A_t) \\ \log(K_t) \end{bmatrix} = \underbrace{\begin{bmatrix} \rho & 0 \\ 1 & \alpha \end{bmatrix}}_{=:M} \begin{bmatrix} \log(A_{t-1}) \\ \log(K_{t-1}) \end{bmatrix} + \begin{bmatrix} \sigma\epsilon_t \\ \log(\alpha\beta) \end{bmatrix} \qquad (11)$$

The two eigenvalues of M are ρ and α, implying that the impact of $\log(A_{t-T})$ and $\log(K_{t-T})$ on $(\log(A_t), \log(K_t))$ decays geometrically at rate $\max\{|\rho|, |\alpha|\}$. This also implies that in order to accurately predict the policy functions based on a truncated history of innovations in a setting where

[19] Since the innovations are i.i.d. by construction, they may be easier for the network to process than highly persistent levels.

ρ is close to 1, the history needs to be very long, leading to high dimensional state $\mathbf{X}_t^{\text{seq},T}$.[20]

3.2.3 Remaining parts of the algorithm

The remaining parts of the algorithm follow Azinovic et al. (2022) (DEQN) and are summarized here for completeness.

Loss function We train the neural network by minimizing a loss function. Specifically, we use the mean squared error in the household optimality condition, (5), expressed in terms of relative consumption error. By minimizing the (weighted) Euler equation errors, we train the neural network policy function to take shape that is consistent with the optimal choices of the representative household.

In order to evaluate the households optimality condition implied by the policy encoded the neural network, $\mathcal{N}_{\boldsymbol{\rho}}(\mathbf{X}_t^{\text{seq},T}) = K_{t+1}$, we need access to the current value of capital and productivity[21] $\mathbf{X}_t^{\text{state}} = [A_t, K_t]$. Hence, even though the neural network's input does only consist of $\mathbf{X}_t^{\text{seq},T}$, we need to also keep track of the associated state vector $\mathbf{X}_t^{\text{state}}$ in order to be able to evaluate the loss function. The distribution over possible $t+1$ histories, $\mathbf{X}_{t+1}^{\text{seq},T}$, is implied by the current history $\mathbf{X}_t^{\text{seq},T}$, together with the stochastic process for the exogenous variable. The distribution over possible $t+1$ states, $\mathbf{X}_{t+1}^{\text{state}} = [A_{t+1}, K_{t+1}] = [A_{t+1}, \mathcal{N}_{\boldsymbol{\rho}}(\mathbf{X}_t^{\text{seq},T})]$, is given by the current state $\mathbf{X}_t^{\text{state}}$, the policy function encoded by the weights of the neural network, and the stochastic processes for the exogenous variables.

We construct the relative Euler equation error by plugging the neural network approximation of the history-based policy function into equation (5).

$$\text{ree}(\mathbf{X}_t^{\text{seq},T}, \mathbf{X}_t^{\text{state}}, \boldsymbol{\rho}) := \frac{(u')^{-1}\left(\mathrm{E}\left[\beta u'(C_{t+1})(1-\delta+r_{t+1}^K)\right]\right)}{C_t} - 1 \tag{12}$$

where

$$C_t = A_t K_t^\alpha + (1-\delta)K_t - K_{t+1} \tag{13}$$

$$C_{t+1} = A_{t+1} K_{t+1}^\alpha + (1-\delta)K_{t+1} - K_{t+2} \tag{14}$$

$$K_{t+1} = \mathcal{N}_{\boldsymbol{\rho}}(\mathbf{X}_t^{\text{seq},T}) \tag{15}$$

$$K_{t+2} = \mathcal{N}_{\boldsymbol{\rho}}(\mathbf{X}_{t+1}^{\text{seq},T}) \tag{16}$$

$$\mathbf{X}_{t+1}^{\text{seq},T} = [A_{t+1-T}, \ldots, A_{t+1}]. \tag{17}$$

We define the loss function as a simple average of the square of the relative Euler equation error

[20] With a slight abuse of notation, we use $\mathbf{X}^{\text{seq},T}$ to indicate the approximate state in sequence space based on either sequences of innovations (ϵ) or values (A).

[21] Because we need to evaluate objects like current period output, or marginal product of capital next period.

evaluated across a dataset $\mathcal{D}$ of state-histories pairs.

$$\mathcal{D} := \{(\mathbf{X}_1^{\text{seq},T}, \mathbf{X}_1^{\text{state}}), (\mathbf{X}_2^{\text{seq},T}, \mathbf{X}_2^{\text{state}}), \ldots, (\mathbf{X}_N^{\text{seq},T}, \mathbf{X}_N^{\text{state}})\} \tag{18}$$

$$\mathcal{L}(\mathcal{D}, \boldsymbol{\rho}) := \frac{1}{N} \sum_{(\mathbf{X}_t^{\text{seq},T}, \mathbf{X}_t^{\text{state}}) \in \mathcal{D}} \text{ree}(\mathbf{X}_t^{\text{seq},T}, \mathbf{X}_t^{\text{state}}, \boldsymbol{\rho})^2. \tag{19}$$

Updating the neural network parameters We update the parameters using the ADAM optimizer (Kingma and Ba, 2014), which is a variant of gradient descent. The update rule implied by vanilla gradient descent is given by

$$\boldsymbol{\rho}^{\text{new}} = \boldsymbol{\rho}^{\text{old}} + \alpha^{\text{learn}} \nabla_{\boldsymbol{\rho}} \mathcal{L}(\mathcal{D}, \boldsymbol{\rho}) \tag{20}$$

The ADAM optimizer modifies the equation (20) by introducing parameter-specific adaptive learning rates rather than using a single constant learning rate for all parameters.

Sampling the most relevant state-history pairs To generate a dataset $\mathcal{D}$ of state-history pairs, we closely follow the ergodic simulation scheme used in the DEQN algorithm. We start the simulation with a random sample of initial exogenous and endogenous states, and then simulate exogenous states using their stochastic process, and evolve endogenous states using the current approximation of the policy functions. Relative to the original DEQN method, our algorithm additionally keeps track of truncated shock histories. This is important because we use the truncated histories of shocks as neural network inputs. A key advantage of using truncated shock histories as network input is that their distribution does not change during the training, as it is pinned-down by model primitives, *i.e.* shock distributions, rather than by policy functions.

3.3 Parameterization

We set relative risk aversion $\gamma = 2$, depreciation $\delta = 0.1$, and patience $\beta = 0.95$. The capital share in production is set to $\alpha = \frac{1}{3}$, and the process for logarithm of TFP has persistence $\rho^A = 0.8$ and $\sigma^A = 0.03$. The model parameters are summarized in table A1.

3.4 Implementation

We train a densely connected feed-forward neural network to predict the savings rate s_t in period t as a function of the last T realizations of innovations to aggregate productivity, *i.e.*, $\mathbf{X}_t^{\text{seq},T} = [\epsilon_{t-T}, \ldots, \epsilon_{t-1}, \epsilon_t] \in \mathbb{R}^T$ and $s_t = \mathcal{N}_{\boldsymbol{\rho}}(\mathbf{X}_t^{\text{seq},T})$. Approximating the savings rate as opposed to directly parameterizing the savings function K_{t+1}, has the advantage that we can ensure that savings rate is bounded in $(0, 1)$ by applying a sigmoid activation in the output layer.[22] Because of that, our policy function is guaranteed to satisfy the budget constraint and also guaranteed to ensure non-negativity of consumption. Given the current capital stock and productivity level, we compute the quantity of resources on hand $M_t = A_t K_t^{\alpha} L^{1-\alpha} + (1-\delta)K_t$. The savings rate and resources on hand then

[22] The sigmoid function is defined as $\text{sigmoid}(x) := \frac{1}{1+e^{-x}}$.

imply the current consumption $C_t = (1 - s_t)M_t$ and the next period's capital $K_{t+1} = s_t M_t$. This illustrates why our algorithm still needs to keep track of values of state variables. Even though states are not used as network input, they are still required to construct objects like resources at hand, marginal products, etc.

We organize the training procedure into a sequence of *episodes*. An episode starts by inheriting a neural policy function and a dataset $\mathcal{D}_j$ consisting of state-history pairs from the previous episode. Then we use the approximate policy function and a pseudo-random number generator to simulate a new training dataset of N^{data} state-history pairs $\mathcal{D}_j = \{(\mathbf{X}_{j,i}^{\text{state}}, \mathbf{X}_{j,i}^{\text{seq},T})\}_{i=1}^{N^{\text{data}}}$:

$$\text{obtain savings rate for old state :} \qquad s_{j,i} = \mathcal{N}_{\boldsymbol{\rho}}(\mathbf{X}_{j,i}^{\text{seq},T}) \tag{21}$$

$$\text{compute implied capital in the next period :} \qquad K_{j+1,i} = s_{j,i}(A_{j,i}K_{j,i}^{\alpha} + (1-\delta)K_{j,i}) \tag{22}$$

$$\text{draw a new random shock :} \qquad \epsilon_{j+1,i} \sim \mathcal{N}(0,1) \tag{23}$$

$$\text{update the truncated history:} \qquad \mathbf{X}_{j+1,i}^{\text{seq},T} = \left[[\mathbf{X}_{j,i}^{\text{seq},T}]_{2:T}, \epsilon_{j+1,i}\right] \tag{24}$$

$$\text{update the agg. state:} \qquad \mathbf{X}_{j+1,i}^{\text{state}} = [K_{j+1,i}, A_{j+1,i}]. \tag{25}$$

Generating new training data in this way is computationally cheap and does not pose a bottleneck to our algorithm. This has the advantage that we can generate new training dataset after every episode, such that no datapoint is used twice during the training. Therefore, overfitting is not a concern. A key difference of our approach, relative to previous deep learning based approaches is that the distribution of neural network inputs, $\mathbf{X}^{\text{seq},T}$, is exogenous and hence remains stationary throughout the training. The state-space representation $\mathbf{X}^{\text{state}}$ on the other hand, is converging to the ergodic set of states as the neural network learns the equilibrium policies.

After obtaining a fresh dataset, we update the policy function by running a short training loop using the Adam optimizer (see Kingma and Ba, 2014) with learning rate α^{learn} and mini-batch size $N^{\text{mini-batch}}$. We evaluate the gradient of the loss function using standard reverse mode automatic differentiation.

3.5 Training

Hyperparameters We use a history of $T = 100$ previous productivity innovations as the approximate sufficient statistic replacing the state vector as a network input. We parameterize the history-based policy function using a densely-connected feed-forward neural network with three hidden layers. Each layer consists of 128 *gelu* activated neurons.[23] The output layer is sigmoid activated, ensuring that the predicted savings rate lies between 0 and 1. To approximate the conditional expectation in the Euler equation we use Gauss-Hermite quadrature with $N^{\text{quad}} = 8$ integration nodes.[24] For minimization of the loss function, we rely on the basic ADAM optimizer with a learning rate of 10^{-5} and mini-batch size of 256. In each episode, our algorithm simulates 4096 state-history pairs, implying 16 gradient descent steps per episode. The hyperparameters are summarized in table A2.

[23] The gelu activation function is given by $f(x) = x\Phi(x)$, where $\Phi(x)$ denotes the Gaussian cumulative distribution function.

[24] An alternative would be to discretize the AR(1) process, as in Rouwenhorst (1995).

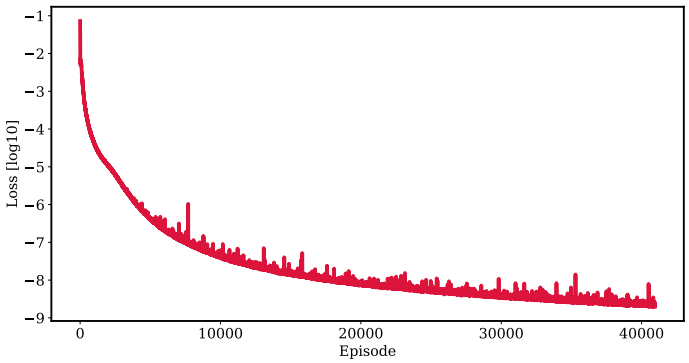


Figure 1: Loss function when training the neural network for ≈40,000 episodes to solve the Brock and Mirman (1972) model. The loss function is the mean squared error in the equilibrium conditions of the model, each episode consists of 4096 simulated states.

Training progress Figure 1 shows the loss function during training. The loss function decreases steadily during the training and reaches a value well below 10^{-8} by the end of training.[25]

3.6 Accuracy

We now turn to assessing the accuracy of our solution. First, we inspect the error in the equilibrium conditions on the simulated ergodic set of states. The single equilibrium condition that characterizes the solution to the model is the Euler equation (5). The left panel in figure 2 shows the distribution of the absolute value of equilibrium condition error expressed in units of relative consumption errors. The vertical lines show the mean, the 99th percentile as well as the 99.9th percentile of the error distribution. As shown in the figure, the errors in the equilibrium conditions are low, with a mean error below 0.005% and the 99.9th percentile of errors below 0.025%.

The error in the equilibrium conditions is an implicit measure, and depending on the model at hand, there might not be an obvious mapping between the magnitude of this implicit error measure and magnitude of actual policy error. For a complicated model, equilibrium conditions error might be the only available error measure. However, in this simple model, we have access to a high-quality reference solution provided by a standard grid-based solution methods. With a dense enough grid, the present model can be solved to very high accuracy, such that we can regard the obtained solution as a good proxy to the *true* policy function. Exploiting this classical solution, in figure 2, we show a comparison of the policy learned by the neural network, to the policy computed with standard policy time iteration (see, *e.g.*, Judd, 1998). The comparison shows two things: first, the difference between the two policies is small, the mean relative difference is below 0.005%, and the 99.9th percentile is below 0.015%. This illustrates that the neural network is able to approximate the equilibrium policies to very high accuracy. Second, when comparing the *explicit* policy error (right panel) to the *implicit* error in the equilibrium conditions (left panel), we can observe that both errors are of the same magnitude. This indicates that the implicit error in the equilibrium conditions is a good proxy for the policy error. Although this is a reassuring finding, it is important to point out

[25]The training run takes roughly 4 minutes on a Tesla T4 GPU.

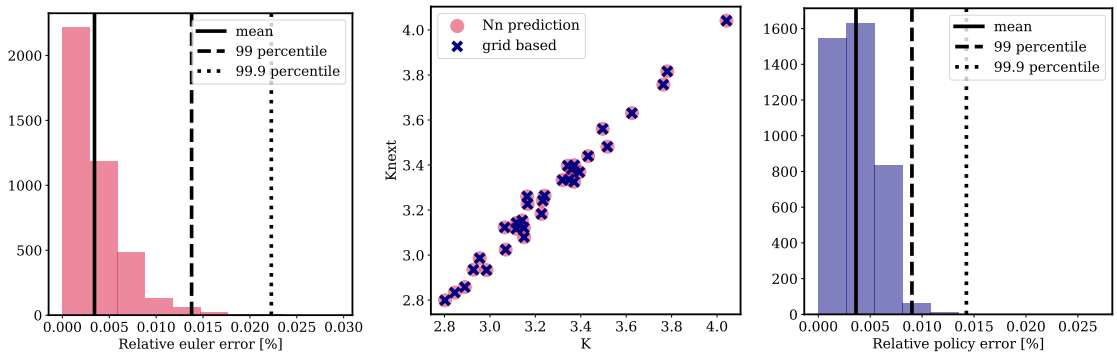


Figure 2: The left panel shows the distribution of errors in the optimality condition (equation (5)) in % on 4096 simulated states after training the neural network. The solid vertical line shows the mean error, the dashed vertical line shows the 99th percentile of errors and the vertical dotted line shows the 99.9th percentile of errors. The middle panel compares the policy learned by the neural network (round dots) to the policy solved for with a conventional grid-based method. The right panel shows the distribution of errors when comparing the policy learned by the neural network to the policy solved with a conventional method.

that this finding depends on the model, so we cannot generalize this conclusion to other models.

4 Application to an OLG model

In this section, we test the performance of our algorithm in a more challenging setup. We compute the approximate equilibrium in an overlapping generations economy featuring portfolio choice, non-trivial market clearing, and more complex stochastic processes for aggregate shocks. In particular, we model 72 overlapping generations of households who consume and save. Households allocate their savings into physical capital and risk-free bonds. While bonds are fully liquid, trading capital is subject to convex portfolio adjustment costs, and return to capital is exposed to aggregate risk. There are three sources of uncertainty: as in the Brock and Mirman (1972) model, the logarithm of TFP follows an AR(1) process. In addition, the economy switches between two regimes: normal times and disaster times. In disaster times, the depreciation to capital is stochastic and on average higher than during normal times. The regime process evolves according to a Markov transition matrix. Since capital is predominantly held by older households, the disaster leads to swings in the intergenerational wealth distribution.

We train the neural network to predict equilibrium prices and policies as a function of truncated histories of aggregate shocks. In this economy, our algorithm keeps track of the last T realizations of each of the aggregate shocks, which include T innovations to the logarithm of TFP, T innovations to the depreciation of capital, and the last T regime realizations. The purpose of this section is to demonstrate that our sequence space method is capable of obtaining highly-accurate approximations to equilibria in economies driven by multivariate stochastic processes. The overlapping generations setup furthermore generates an obvious long-range dependence of current outcomes on the realized sequence of past shocks.

4.1 Model

4.1.1 Technology

A representative firm operates a Cobb-Douglas technology and produces a perishable consumption good using labor and capital

$$Y_t = A_t K_t^{\alpha} L_t^{1-\alpha}. \tag{26}$$

A_t denotes the stochastic total factor productivity (TFP). The logarithm of TFP evolves according to an AR(1) process

$$\log(A_t) = \rho \log(A_{t-1}) + \sigma^A \epsilon_t^A,\ \epsilon_t^A \sim \mathcal{N}(0,1). \tag{27}$$

As in the stochastic growth economy of Brock and Mirman (1972), the firm faces competitive input and output markets. We choose the consumption good as the numéraire, hence price of consumption is normalized to unity. We denote the wage by w_t and rental rate of capital by r_t^K.

4.1.2 Demographics

Each period a new representative cohort of mass 1 enters the economy and lives for H periods. We denote the age of a cohort by $h \in \{1, \ldots, H\}$. Households are endowed with l^h efficiency units of labor, which they supply inelastically at equilibrium wage. The labor endowment is normalized such that $L = \sum_{h=1}^{H} l^h = 1$. The age dependence of efficient labor supply serves as a device to model a life-cycle profile of household labor income. We abstract from bequest motives, taxes and social security.

4.1.3 Preferences

Households have time separable expected utility over consumption streams. We assume that the felicity function takes the constant relative risk aversion form with coefficient γ. Hence, households rank stochastic consumption stream according to

$$\mathrm{E}\left[\sum_{i=0}^{H-h} \beta^i u(c_{t+i}^{h+i})\right], \tag{28}$$

where β denotes their patience.

4.1.4 Asset markets

Households are able to re-allocate their consumption across time and states by trading in two assets. Two assets and infinitely many possible shock realizations next period[26] imply an incomplete-markets environment. Furthermore, there are two additional financial frictions. First, households are subject to short sale limits on capital and bonds. Second, while the bond is a liquid asset, capital is an

[26] We assume innovations to the logarithm of TFP and capital depreciation to be Gaussian.

illiquid asset and households have to pay adjustment costs when adjusting their capital holdings. The adjustment costs are given by $\psi(k_{t+1}^{h+1}, k_t^h) = \xi^{\text{adj}} \left(k_{t+1}^{h+1} - k_t^h\right)^2$. Short sale constraints on both assets are specified as exogenous limits $b_{t+1}^{h+1} \geq \underline{b}$ and $k_{t+1}^{h+1} \geq \underline{k}$. The net supply of bonds is fixed at B. Bonds can be purchased or sold at equilibrium price p_t, and they deliver a risk-free payoff of 1 in period $t+1$. Purchasing a unit of capital in period t promises a risky payout of $1 - \delta_{t+1} + r_{t+1}^K$ in period $t+1$. There are two sources of risk in returns to capital: first, the rental rate of capital r_{t+1}^K depends on the total factor productivity A_{t+1} in period $t+1$. Second, the depreciation rate of capital in period $t+1$ is stochastic as well and given by

$$\delta_{t+1} = \delta \frac{2}{1 + D_{t+1}}, \text{ where } \log(D_{t+1}) = \begin{cases} \rho^\delta \log(D_t) & \text{normal times in } t+1 \\ \mu^\delta + \rho^\delta \log(D_t) + \sigma^\delta \epsilon_{t+1}^\delta & \text{disaster in } t+1 \end{cases} \tag{29}$$

with $\mu^\delta < 0$ and $\epsilon_{t+1}^\delta \sim \mathcal{N}(0,1)$. Normal and disaster times are modeled as two discrete Markov regimes with transition matrix $\Pi^{\text{regime}} = \begin{bmatrix} \pi^{n \to n} & 1 - \pi^{n \to n} \\ 1 - \pi^{d \to d} & \pi^{d \to d} \end{bmatrix}$, where $\pi^{n \to n}$ and $\pi^{d \to d}$ denote the persistence of the normal and disaster regimes, respectively.

4.1.5 Household problem

Recursively formulated, the households' problem is characterized by the following Bellman equation

$$V_t^h = \max_{k_{t+1}^{h+1}, b_{t+1}^{h+1}} u(c_t^h) + \beta \mathrm{E}\left[V_{t+1}^{h+1}\right] \tag{30}$$

subject to :

$$\begin{aligned} c_t^h = & \underbrace{l^h w_t}_{\text{lab. inc.}} + \underbrace{b_t^h + (1 - \delta_t + r_t^K) k_t^h}_{\text{payout of assets}} - \underbrace{(p_t b_{t+1}^{h+1} + k_{t+1}^{h+1})}_{\text{savings}} - \underbrace{\psi(k_{t+1}^{h+1}, k_t^h)}_{\text{adj. costs}} \\ & \underline{b} \leq b_{t+1}^{h+1} \\ & \underline{k} \leq k_{t+1}^{h+1}. \end{aligned}$$

In every period, households receive labor income and the payoff of their asset holdings, and decide how much to consume and how much to invest in each of the two assets, subject to adjustment costs and short-sale constraints. The time index denotes the dependence on the aggregate state variables.

The solution to the household problem is characterized by a set of Karush-Kuhn-Tucker (KKT)

conditions. For each age-group $h = 1, \ldots, H-1$ and each asset, we have two sets of KKT conditions

$$p_t u'(c_t^h) = \beta \mathrm{E}\left[u'(c_{t+1}^{h+1})\right] + \lambda_t^{b,h} \tag{31}$$

$$0 = \lambda_t^{b,h}(b_{t+1}^{h+1} - \underline{b}) \tag{32}$$

$$0 \leq (b_{t+1}^{h+1} - \underline{b}) \tag{33}$$

$$0 \leq \lambda_t^{b,h} \tag{34}$$

$$\left(1 + \frac{\partial}{\partial k_{t+1}^{h+1}}\psi(k_{t+1}^{h+1}, k_t^h)\right) u'(c_t^h) = \beta \mathrm{E}\left[u'(c_{t+1}^{h+1})\left(1 - \delta_{t+1} + r_{t+1}^k - \frac{\partial}{\partial k_{t+1}^{h+1}}\psi(k_{t+2}^{h+2}, k_{t+1}^{h+1})\right)\right] + \lambda_t^{k,h} \tag{35}$$

$$0 = \lambda_t^{k,h}(k_{t+1}^{h+1} - \underline{k}) \tag{36}$$

$$0 \leq (k_{t+1}^{h+1} - \underline{k}) \tag{37}$$

$$0 \leq \lambda_t^{k,h}. \tag{38}$$

For each age-group and each asset we use the Fischer-Burmeister function (see, Jiang, 1999) $\psi^{FB}(a,b) := a - b - \sqrt{a^2 + b^2}$ to collapse the set of KKT conditions into a single equation. The above KKT conditions are fully characterized by

$$0 = \psi^{FB}\left(p_t u'(c_t^h) - \beta \mathrm{E}\left[u'(c_{t+1}^{h+1})\right], b_{t+1}^{h+1} - \underline{b}\right) \tag{39}$$

$$0 = \psi^{FB}\left(\left(1 + \frac{\partial}{\partial k_{t+1}^{h+1}}\psi(k_{t+1}^{h+1}, k_t^h)\right) u'(c_t^h) - \beta \mathrm{E}\left[u'(c_{t+1}^{h+1})\left(1 - \delta_{t+1} + r_{t+1}^k - \frac{\partial}{\partial k_{t+1}^{h+1}}\psi(k_{t+2}^{h+2}, k_{t+1}^{h+1})\right)\right], (k_{t+1}^{h+1} - \underline{k})\right) \tag{40}$$

Analogously to equation (12) for the representative agent model, we rearrange the equations so that the deviations from zero are expressed in units of relative consumption error.

4.1.6 Equilibrium

Now, we define the recursive equilibrium of the economy

State space The state of the economy

$$\mathbf{X}_t^{\text{state}} = [\chi_t, A_t, \delta_t, \mu_t] \in \{0,1\} \times \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^{H-1+H-2} \tag{41}$$

is given by the regime indicator, χ_t, total factor productivity A_t, capital depreciation δ_t, and the distribution of assets across the age groups $\mu_t = \{(b_t^h, k_t^h)\}_{h=1}^H$.[27]

[27] Strictly speaking, because we assume that households enter the economy without assets, and because of market clearing, a sufficient statistic for the asset distribution consists of $H-1$ free parameters for capital and $H-2$ for the bond, because the bond supply is fixed.

Functional rational expectations equilibrium The functional rational expectations equilibrium of our economy consists of H policy functions $\pi^{k,h}(\mathbf{X}_t^{\text{state}}) = k_{t+1}^{h+1}$ for investment in capital, H policy functions $\pi^{b,h}(\mathbf{X}_t^{\text{state}}) = b_{t+1}^{h+1}$ for investment in the bond, and a price function $\pi^{p}(\mathbf{X}_t^{\text{state}}) = p_t$, such that the KKT conditions hold and the bond and capital markets clear.[28] The aim of our algorithm is to compute approximations to these equilibrium functions.

4.2 Parameterization

We set the model parameters to standard values in the literature. One model period corresponds to one year, and we model $H = 72$ overlapping generations. The model ages $h = 1$ to 72 correspond to adult life from 21 years to 92 years. The age dependent efficient units of labor are chosen to generate a life-cycle profile of labor income and are shown in figure A1. We set the coefficient of relative risk aversion at $\gamma = 2$ and patience to $\beta = 0.96$. We assume strict borrowing constraints with $\underline{b} = \underline{k} = 0$. The adjustment costs for capital are set to $\xi^{\text{adj}} = 0.1$. The persistence of the logarithm of TFP is set to $\rho = 0.85$, and the standard deviation of innovations to the logarithm of TFP is set to $\sigma^A = 0.03$. The long-run depreciation of capital in normal times is given by $\delta = 0.1$. The parameters governing the depreciation of capital in the disaster regime are given by $\rho^\delta = 0$, $\sigma^\delta = 0.2$ and $\mu^\delta = -1.10$, implying 50% higher depreciation in the disaster when $\log(D_t) = \mu^\delta$. The per-period probability to enter into the disaster state is 6% and the per-period probability to exit the disaster state is 33.33%, implying $\pi^{n\to n} = 0.94$ and $\pi^{d\to d} = 2/3$. The parameter values are summarized in Table A3.

4.3 Implementation

We use a densely-connected feed-forward neural network to parameterize a key set of functions and then obtain the remaining equilibrium objects in closed form. These key functions consist of the bond price, p_t, as well as the investment in capital, k_{t+1}^{h+1}, and investment in the bond, b_{t+1}^{h+1}, for each age-group.[29] We use the market clearing layers approach (as in Azinovic and Zemlicka, 2024), so the neural network architecture ensures that the predicted policies are always consistent with bond market clearing.

Following our sequence space approach, the neural network predicts the policy and price functions as a function of truncated history of aggregate shocks. Hence, the input to the neural network is given by

$$\mathbf{X}_t^{\text{seq},T} = [\underbrace{\chi_{t-(T-1)}, \ldots, \chi_{t-1}, \chi_t}_{T \text{ last regime indicators}}, \underbrace{\epsilon^A_{t-(T-1)}, \ldots, \epsilon^A_{t-1}, \epsilon^A_t}_{T \text{ last innovations to TFP}}, \underbrace{\epsilon^\delta_{t-(T-1)}, \ldots, \epsilon^\delta_{t-1}, \epsilon^\delta_t}_{T \text{ last innovations to deprec.}}] \in \mathbb{R}^{3T}. \tag{42}$$

As before, we train the neural network to minimize the errors in all equilibrium conditions. The resulting loss function consists of the $2 \times (H-1)$ optimality conditions for the capital and bond

[28] We use the first-order conditions of representative firm directly to compute the equilibrium wage and rental rate as a function of the state of the economy. The goods market clears by Walras' law.

[29] To be precise, we only need to predict the asset choice for $H-1$ age groups, since, for our calibration, it is always optimal to consume everything in the last period of life. Similarly to our approach in the Brock and Mirman (1972) model, the neural network predicts the share of cash-at-hand invested in each of the two assets, which we then use to construct the prediction for the asset choices.

policies for each age group, except the last who consumes everything. As in the previous example, the training data in episode j are given by state-history pairs $\mathcal{D}_j = \{(\mathbf{X}_{j,i}^{\text{state}}, \mathbf{X}_{j,i}^{\text{seq},T})\}_{i=1}^{N^{\text{data}}}$ and use the Adam optimizer (see Kingma and Ba, 2014). After each training episode we simulate states and histories forward to generate a new training dataset $\mathcal{D}_{j+1}$. We approximate the expectations over the continuous shocks ϵ_{t+1}^A and ϵ_{t+1}^δ using Gauss-Hermite quadrature.

4.4 Training

Our training procedure follows the stepwise approach to solve models with multiple assets outlined in Azinovic and Zemlicka (2024). We proceed along four steps:

First, we set the net supply of bonds to zero. Given a strict no-short sale constraint the market clearing condition implies that all equilibrium bond holdings are zero. Hence we can first restrict our attention to solving for capital savings policy functions while keeping bond savings fixed at zero.

Second, we train the neural network price function to learn the bond price implied by consumption allocations of the capital-only economy. We plug the set of consumption functions obtained by solving the capital-only economy into bond Euler equations. We invert the Euler equations to recover implied bond prices. The largest of the prices would be the price of the first ϵ of bond supply introduced into this economy. We train the price network to replicate this price as a function of truncated histories of aggregate shocks.

Third, the previous two steps provide us with a good initial guess for the policy and price functions in an economy with a comparatively small bond supply. We now proceed to slowly increase the bond supply. After every small increase in the bond supply, we retrain the neural network using the previous solution as a starting point. We proceed until the full bond supply is reached.

Fourth, after reaching the full supply, we continue training the neural network on the final model specifications until we reach desired level of accuracy. The hyperparameters for our implementations are summarized in table A4.

4.5 Accuracy

We assess the accuracy of our neural network solution by investigating the errors in the equilibrium conditions over the simulated ergodic set. Figure 3 shows the errors in the equilibrium conditions for each age group along the simulated paths of the economy. The errors are expressed in units of relative consumption errors. The model is solved accurately, the mean and 90th percentile of errors are below 0.1% and the 99.9th percentile below 0.32% for both assets and all the age groups. The mean error is around 0.03%.

In the middle panel in figure 3 we show the distribution of asset holdings over the age groups. The savings behavior of households is driven by the life-cycle forces. Facing an increasing labor income profile, young households are borrowing constrained. Later in life, they start to accumulate savings to make sure they are able to finance their later-age consumption after the drop in their labor endowment after the age of 62.[30]

[30] While we abstract from modeling retirement explicitly, the reduction in labor efficiency units around the age of 60 serves as a proxy for the income decline associated with retirement.

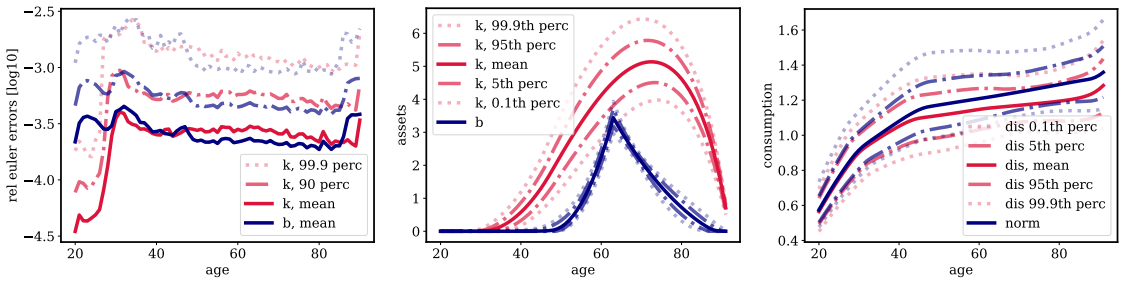


Figure 3: The left panel shows the final level of errors in the equilibrium conditions achieved by our training procedure. The dotted line shows the 99.9th percentile, the dash-dotted line the 90th percentile, and the solid line shows the mean. The red lines show the errors in the optimality conditions for capital for each age group, and the blue lines show the errors in the optimality conditions for the bonds. The middle panel shows the distribution of assets over the simulated ergodic set of states. The right panel shows the resulting consumption by age group. The blue lines in the right panel show consumption conditional on the economy being in the normal state, and the red lines show consumption statistics conditional on the economy being in the disaster state.

The right panel in figure 3 shows the distribution of consumption across the age groups, conditional on the regime of the economy. The disaster, modeled as an increase in the capital depreciation and uncertainty, leads to the largest consumption decline for households around the age 80.

5 Application to heterogeneous households and heterogeneous firms

In this section, we apply our algorithm to compute an approximate equilibrium in an economy with heterogeneous households and heterogeneous firms. The households and the firms are both subject to uninsurable idiosyncratic risk, as well as to aggregate productivity shocks, and swings in the level of aggregate and idiosyncratic volatility. Firms operate a decreasing returns to scale technology that produces the final good using capital and labor. The firms own their capital, hire workers on a Walrasian labor market, and pay dividends to their shareholders. Firms make their decisions to maximize the present discounted value of their dividends. We assume that firms use a weighted average of households' marginal rates of substitution to discount future cash flows. The weighting is proportional to households' stock holdings. Households supply labor inelastically at equilibrium wage, consume, and invest in the aggregate stock market index. Besides the fluctuations in wages and stock prices induced by aggregate shocks, households are subject to uninsurable idiosyncratic labor endowment shocks. The state space of the economy includes two endogenously moving distributions: the firm distribution over capital and idiosyncratic firm productivity, as well as the household distribution over stock holdings and idiosyncratic labor productivity. The distribution of firms is relevant to households because it determines the dividends paid on stock holdings, and the wage rate on supplied labor. In turn, the stochastic discount factor used by firms to make their investment decisions depends on the distribution of households.

5.1 Model

5.1.1 Firms and technology

Time is discrete and infinite, $t = 0, 1, \ldots$. There is a continuum of firms that face idiosyncratic and aggregate risk featuring stochastic volatility in the spirit of Bloom et al. (2018). At the beginning of each period, there is a continuum of firms of mass 1. We assume that only a fraction Γ of firms survive to the end of the period and produce the consumption good. A fraction $(1-\Gamma)$ of firms exits the economy, and is replaced by an equal measure of start-ups that, conditional on surviving, start to produce in the next period.[31] Each of the firms $i \in [0, \Gamma]$ operates a decreasing returns to scale production technology

$$y_t^i = A_t z_t^i (k_t^i)^\alpha (l_t^i)^{1-\alpha-\zeta}, \tag{43}$$

where A_t denotes stochastic aggregate productivity, which affects all firms in the economy, z_t^i denotes the realization of firm specific productivity. The parameter α determines the elasticity of the output with respect to capital, and the parameter ζ controls the degree of returns to scale.[32] Aggregate productivity follows an AR(1) process

$$\log(A_t) = \rho^A \log(A_{t-1}) + \sigma_{t-1}^A \epsilon_t^A, \tag{44}$$

where $\epsilon_t^A \sim \mathcal{N}(0,1)$. The volatility of innovations to aggregate TFP, $\sigma_t^A = \sigma(U_t) \in \{\sigma_L^A, \sigma_H^A\}$, takes two values depending on the uncertainty regime $U_t \in \{L, H\}$. The uncertainty regime follows a Markov chain as in Bloom et al. (2018). We denote the transition probabilities between the uncertainty regimes with $\pi_{L,H}^U$, $\pi_{L,L}^U = 1 - \pi_{L,H}^U$, $\pi_{H,L}^U$, and $\pi_{H,H}^U = 1 - \pi_{H,L}^U$, respectively. The firm-level idiosyncratic productivity follows a three-state Markov process $z_t \in \{z_1, z_2, z_3\}$, with $z_2 = (z_1 + z_3)/2$. The transition probabilities vary with the uncertainty regime: $\Pi_t^z \in \{\Pi_L^z, \Pi_H^z\}$. We are interested in the effect of a pure uncertainty shock, *i.e.* in an increase in the variance of future idiosyncratic productivity without changing the current productivity or the expected future productivity level $\mathrm{E}\left[z_{t+1}^i | z_t\right]$. To do so, assume that the transition matrices in the low and high uncertainty regimes are given by

$$\Pi_L^z = \begin{bmatrix} \pi_{11}^z & \pi_{12}^z & \pi_{13}^z \\ \pi_{21}^z & \pi_{22}^z & \pi_{23}^z \\ \pi_{31}^z & \pi_{32}^z & \pi_{33}^z \end{bmatrix}, \Pi_H^z = \begin{bmatrix} \pi_{11}^z + u & \pi_{12}^z - 2u & \pi_{13}^z + u \\ \pi_{21}^z + u & \pi_{22}^z - 2u & \pi_{23}^z + u \\ \pi_{31}^z + u & \pi_{32}^z - 2u & \pi_{33}^z + u \end{bmatrix}, \tag{45}$$

where $0 \leq u < \min\{\frac{1}{2}\pi_{12}^z, \frac{1}{2}\pi_{22}^z, \frac{1}{2}\pi_{13}^z\}$ parameterizes the increase in the idiosyncratic uncertainty. This formulation allows us to keep the idiosyncratic productivity level z_t^i, as well as the expected future productivity $\mathrm{E}\left[z_{t+1}^i | z_t\right]$ constant, when the economy transitions into the high volatility

[31] Firm exit is modeled as in Azinovic et al. (2023). We assume that the capital of exiting firms is destroyed in the process.

[32] The capital share is $\alpha/(1-\zeta)$.

regime.[33] Firms accumulate capital subject to capital adjustment costs of the form

$$\psi(k', k) = \frac{\xi(k', k)}{2} k \left(\frac{k'}{k} - 1 \right)^2, \tag{46}$$

where

$$\xi(k', k) = \begin{cases} \xi^{\text{up}} & \text{for } k' > k \\ \xi^{\text{down}} & \text{for } k' \leq k. \end{cases} \tag{47}$$

We assume that $\xi^{\text{down}} > \xi^{\text{up}}$, so reducing the capital stock is associated with larger adjustment costs than increasing the capital stock by the same percentage. Although the adjustment costs parameter changes discontinuously at $k' = k$, $\psi(k', k)$ is not only continuous but also differentiable. At $k' = k$ we have $\psi(k', k) = \frac{\partial \psi(k', k)}{\partial k'} = \frac{\partial \psi(k', k)}{\partial k} = 0$.[34]

Each firm sets its investment and dividend payout policies to maximize its objective function: the present discounted value of its dividends, subject to a budget constraint, and a non-negative dividends constraint. The firm problem is characterized by the following Bellman equation:

$$v_t(z_t^j, k_t^j) = \max_{k_{t+1}^j \geq 0} d_t^j + \Gamma \, \mathrm{E}_t \left[\Lambda_{t+1,t} \, v_{t+1}(z_{t+1}^j, k_{t+1}^j) \right] \tag{48}$$

subject to:

$$d_t^j = y_t^j - w_t l_t^j - i_t^j - \psi(k_{t+1}^j, k_t^j) \tag{49}$$

$$d_t^j > \underline{d}. \tag{50}$$

Γ denotes the mass of surviving firms, which is equal to the (idiosyncratically independent) survival probability of an individual firm between t and $t+1$. The time index of the value function stands for the dependence of the firm problem on the aggregate state of the economy. $\Lambda_{t+1,t}$ denotes the aggregate stochastic discount factor given by the distribution of shareholders. The policy functions

[33] To the best of our knowledge, maintaining these two features would not easily be possible using, for example, a Rouwenhorst (1995) or Tauchen (1986) discretization methods for an AR(1) process.

[34] $\xi(k', k)$ is not continuous at $k' = k$, but the quadratic term $(k'/k - 1)^2$ ensures that $\psi(k', k)$ remains differentiable. Additionally, we smooth out the jump in the adjustment cost parameter such that $\xi(k', k) = w^{\text{up}} \xi^{\text{up}} + (1 - w^{\text{up}}) \xi^{\text{down}}$, where $w^{\text{up}} = \text{sigmoid}\left(s \frac{k' - k}{k} \right)$, and where $s > 0$ denotes a smoothing parameter. For large s, w^{up} converges to a step function, *i.e.* $w^{\text{up}} \approx 1$ for $s \frac{k' - k}{k} >> 0$ and $w^{\text{up}} \approx 0$ for $s \frac{k' - k}{k} << 0$.

of the firms are characterized by a set of KKT conditions.

$$\left(1+\frac{\partial}{\partial k_{t+1}^i}\psi(k_{t+1}^i,k_t^i)\right)(1+\lambda_t^i)=\Gamma\mathrm{E}\left[\Lambda_{t+1,t}\left(1+r_{t+1}^{i,k}-\frac{\partial}{\partial k_{t+1}^i}\psi(k_{t+2}^i,k_{t+1}^i)\right)(1+\lambda_{t+1}^i)\right]$$

$$\Leftrightarrow \lambda_t^i=\frac{\Gamma\mathrm{E}\left[\Lambda_{t+1,t}\left(1+r_{t+1}^{i,k}-\frac{\partial}{\partial k_{t+1}^i}\psi(k_{t+2}^i,k_{t+1}^i)\right)(1+\lambda_{t+1}^i)\right]}{\left(1+\frac{\partial}{\partial k_{t+1}^i}\psi(k_{t+1}^i,k_t^i)\right)}-1 \tag{51}$$

$$0\leq(d_t^i-\underline{\mathrm{d}}) \tag{52}$$

$$0\leq\lambda_t^i \tag{53}$$

$$0=(d_t^i-\underline{\mathrm{d}})\lambda_t^i, \tag{54}$$

where λ_t^i denotes the Lagrange multiplier on the non-negative dividend constraint and where $r_t^{i,k}$ denotes the marginal product of capital installed in the firm

$$r_t^{i,k}=A_t z_t^i \alpha (k_t^i)^{\alpha-1}(l_t^i)^{1-\alpha-\zeta}-\delta. \tag{55}$$

Using the Fischer-Burmeister function, we rewrite the KKT conditions of the firm problem as a single equation $\psi^{FB}(\lambda_t^i, d_t^i-\underline{\mathrm{d}})=0$, where λ_t^i is given in equation (51).

Aggregation of firms Firms hire efficiency units of labor on a competitive spot market. There is a single equilibrium wage w_t per efficiency unit of labor, which is paid by all the firms. The labor demand of an individual firm is given by a static first-order condition

$$w_t=(1-\alpha-\zeta)\frac{y_t^i}{l_t^i}=(1-\alpha-\zeta)A_t z_t^i (k_t^i)^{\alpha}(l_t^i)^{-\alpha-\zeta} \tag{56}$$

$$\Leftrightarrow l_t^i=\left(A_t z_t^i \frac{1-\alpha-\zeta}{w_t}(k_t^i)^{\alpha}\right)^{\frac{1}{\alpha+\zeta}}. \tag{57}$$

Let $\mu_t^F(z^i,k^i)$ denote the distribution of surviving firms over idiosyncratic productivity and capital. The overall mass of surviving firms is given by $\Gamma=\int_i \mu_t^F(z^i,k^i)di$. Since we assume that all households supply their labor inelastically, the aggregate labor supply is constant at $L=1$. This allows us to solve for the equilibrium wage as a function of the productivity level A_t, and the cross-sectional distribution of firms. We obtain

$$w_t=A_t(1-\alpha-\zeta)\left(\int\left(z^i(k^i)^{\alpha}\right)^{\frac{1}{\alpha+\zeta}}\mu_t^F(z^i,k^i)di\right)^{\alpha+\zeta}. \tag{58}$$

Given the expression for the wage (58), the labor demand is given in a closed form, as described in equation (57). Note that, because of decreasing returns to scale, the full firm distribution matters for the determination of the aggregate wage. Similarly, aggregate dividends are given by

$$D_t=\int d_t(z^i,k^i)\mu_t(z^i,k^i)di. \tag{59}$$

Startups In each period, households trade shares in the aggregate index of existing firms. Owning a share entitles a household to receive a fraction of the aggregate dividend stream. The firms that exist in period t include a measure Γ of surviving firms that produce in period t, as well as a measure $(1-\Gamma)$ of period t startups, which start producing only in period $t+1$. Holding θ_t^i shares, which have been purchased in period $t-1$, entitles households to a payout of $\theta_t^i(\Gamma p_t + D_t)$. This is because period t startups can only be traded once they exist. The stocks traded in period $t-1$ therefore only make up fraction Γ of the stocks traded in period t.

A measure $(1-\Gamma)$ of startups replaces exiting firms every period. The average value of period t startups is $(1-\Gamma)p_t$. We set the initial capital of startups to be equal to the capital of the firms they are replacing. The required investment is provided uniformly by all the households who therefore also own the startups. The initial investment is not subject to adjustment costs. The size of startup investment per household is given by

$$I_t^{\text{su}} = (1-\Gamma)\int k_{t+1}^i(z_t^i, k_t^i)\frac{\mu(z_t^i, k_t^i)}{\Gamma}\mathrm{d}i. \tag{60}$$

The aggregate rents from startup creation are given by

$$\Pi_t^{\text{su}} = (1-\Gamma)p_t - I_t^{\text{su}}. \tag{61}$$

The rents from startup creation are equal to the difference between the average capital stock in the economy and the price of an equity share. We assume that the profits from startup creation are equally distributed to all households.

5.1.2 Households

There is a unit-mass continuum of infinitely lived households. Households have time separable expected utility preferences over stochastic consumption streams

$$U\left(\{c_t\}_{t=0}^{\infty}\right) = \mathrm{E}\left[\sum_{t=0}^{\infty}\beta^t u(c_t)\right], \text{ where } u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \tag{62}$$

and where β denotes the patience parameter, and γ denotes the coefficient of relative risk aversion. Households' labor endowment is stochastic and follows an AR(1) process

$$\log(e_t^i) = \rho^e \log(e_{t-1}^i) + \sigma^e \epsilon_t^{e,i}, \text{ where } \epsilon_t^{e,i} \sim \mathcal{N}(0,1). \tag{63}$$

We discretize the idiosyncratic labor endowment process into a two-state Markov process using the Rouwenhorst (1995) algorithm. Households supply their efficient units of labor inelastically at equilibrium wage. While households can save by purchasing equity shares θ_t^i, their borrowing is subject to a short-sale limit $\theta_{t+1}^i \geq \underline{\theta}$. The budget constraint of household i reads as

$$c_t^i = e_t^i w_t + \theta_t^i(D_t + \Gamma p_t) - \theta_{t+1}^i p_t + \Pi_t^{\text{SU}}. \tag{64}$$

The solution to the household savings problem is characterized by the following set of KKT conditions

$$p_t u'(c_t^i) = \beta \mathrm{E}\left[u'(c_{t+1}^i)(D_{t+1} + \Gamma p_{t+1})\right] + \lambda_t^i \tag{65}$$

$$0 \leq \lambda_t^i \tag{66}$$

$$0 \leq (\theta_{t+1}^i - \underline{\theta}) \tag{67}$$

$$0 = (\theta_{t+1}^i - \underline{\theta})\lambda_t^i. \tag{68}$$

Households are exposed to idiosyncratic risk because of their labor endowments and to aggregate risk via the wage, dividends, the stock price, and the profits from startup creation. Because of that, aggregate shocks lead to fluctuations in the wealth distribution. We denote the household distribution by $\mu_t^H(e^i, \theta^i)$, where e denotes idiosyncratic labor endowment, and θ share holdings.

5.1.3 Asset markets

Households trade claims on the aggregate dividend stream of *existing* firms. The aggregate asset demand of households is given by

$$\Theta_{t+1}^H = \int \theta_{t+1}^i(e^i, \theta^i)\mu^H(e^i, \theta^i)di. \tag{69}$$

Market clearing requires that $\Theta_{t+1}^H = 1$. Firms take the stochastic discount factor as given. We assume that firms use the intertemporal marginal rate of substitution of shareholder households weighted by their corresponding share holdings as the stochastic discount factor. For a specific realization of aggregate shocks in period $t+1$ and household i with idiosyncratic labor endowment e^i and shares θ^i, we define its intertemporal marginal rate of substitution as

$$\Lambda_{t+1,t}^i(e^i, \theta^i) = \frac{\beta \mathrm{E}_{e_{t+1}^i}\left[u'(c_{t+1}^i)\right]}{u'(c_t^i)}, \tag{70}$$

where the expectations operator is taken over the realizations of idiosyncratic labor endowment. The aggregate stochastic discount factor is given by

$$\Lambda_{t+1,t} = \frac{\int \Lambda_{t+1,t}^i(e^i, \theta^i)\theta_{t+1}(e^i, \theta^i)\mu_t^H(e^i, \theta^i)di}{\int \theta_{t+1}(e^i, \theta^i)\mu_t^H(e^i, \theta^i)di}. \tag{71}$$

5.1.4 Equilibrium

State space The aggregate state of the economy is given by

$$\mathbf{X}_t^{\text{agg. state}} = [\chi_t, A_t, \mu_t^F, \mu_t^H] \in \{0, 1\} \times \mathbb{R}_+ \times \mathbb{R}_+^{N_z \times N_k} \times \mathbb{R}_+^{N_e \times N_\theta}, \tag{72}$$

where χ_t is an indicator denoting the uncertainty regime, A_t denotes the aggregate productivity level, and μ_t^H and μ_t^F denote the distribution of households and firms. We represent the distributions using a finite histogram approximation (see Young, 2010) with $N_z \times N_k$ and $N_e \times N_\theta$ histogram points for firms and households. N_z and N_e denote the number of histogram points along the exogenous

idiosyncratic productivity dimension and N_k and N_θ denote the number of histogram points for the capital and asset holdings respectively.

Functional rational expectations equilibrium The functional rational expectations equilibrium consists of firm policy functions, $\pi^F(\mathbf{X}_t^{\text{agg. state}}, z_t^i, k_t^i) = k_{t+1}^i$, household policy function, $\pi^H(\mathbf{X}_t^{\text{agg. state}}, e_t^i, \theta_t^i) = \theta_{t+1}^i$, as well as a price function $\pi^q(\mathbf{X}_t^{\text{agg. state}}) = p_t$, consistent with optimizing households, optimizing firms, and market clearing.

5.2 Parameterization

One model period corresponds to one calendar year. We set patience to $\beta = 0.95$ and the coefficient of relative risk aversion to $\gamma = 2$. For the idiosyncratic labor process we choose $\rho^e = 0.871$ and $\sigma^e = 0.246$, as in Auclert et al. (2021).[35] We use the Rouwenhorst (1995) method to discretize the labor process into a two state Markov chain with states $e_0 = 0.538, e_1 = 1.462$ and transition matrix

$$\pi^e = \begin{bmatrix} 0.935 & 0.065 \\ 0.065 & 0.935 \end{bmatrix} \tag{73}$$

The short-sale limit is set to $\underline{\theta} = 0$.

Following Bloom et al. (2018), we set the depreciation rate of capital to $\delta = 0.1$ and we choose $\alpha = 0.25$ and $\zeta = 0.25$, consistent with a capital share of $1/3$ in production. The persistence of the logarithm of aggregate productivity is set to $\rho^A = 0.8145$ and the standard deviations of innovations to productivity are set to $\sigma_L^A = 0.0124$ and $\sigma_H^A = 0.0199$, consistent with the quarterly process in Bloom et al. (2018) and Khan and Thomas (2008). We set the survival rate of firms to $\Gamma = 0.965$. There are $N_z = 3$ idiosyncratic productivity levels for firms corresponding to $z_0 = 0.5$, $z_1 = 1.0$, and $z_2 = 1.5$. The corresponding Markov transition matrices are given by

$$\Pi_L^z = \begin{bmatrix} 0.850 & 0.150 & 0.000 \\ 0.075 & 0.850 & 0.075 \\ 0.000 & 0.150 & 0.850 \end{bmatrix} \text{ and } \Pi_H^z = \begin{bmatrix} 0.925 & 0.000 & 0.075 \\ 0.150 & 0.700 & 0.150 \\ 0.075 & 0.000 & 0.925 \end{bmatrix}. \tag{74}$$

The transition matrix between the two uncertainty regimes is given by

$$\Pi^U = \begin{bmatrix} 0.90 & 0.10 \\ 0.21 & 0.79 \end{bmatrix}, \tag{75}$$

implying a quarterly persistence of the high uncertainty regime of 0.943 and a quarterly persistence of the low uncertainty regime of 0.974, as in Bloom et al. (2018). The adjustment costs parameters are set to $\xi^{\text{up}} = 1.0$ and $\xi^{\text{down}} = 2.5$.[36] The parameter values are summarized in table A5.

[35] We convert the quarterly process in Auclert et al. (2021) into a yearly process.

[36] The smoothing parameter in the adjustment cost function is set to $s = 400$, such that $\xi(1.01k, k) = 0.993\xi^{\text{up}} + 0.007\xi^{\text{down}}$ and $\xi(0.99k, k) = 0.007\xi^{\text{up}} + 0.993\xi^{\text{down}}$.

5.3 Implementation

As in the previous examples, we solve the model by training neural networks to approximate a core set of equilibrium functions, and then solve for the remaining equilibrium objects in closed form. Relative to the previous two applications, this model poses an additional challenge: individual policy functions of firms and households depend not only on aggregate quantities, but also on their corresponding idiosyncratic states. To solve for the firm and household policies, we train neural networks to approximate a mapping from truncated histories of aggregate shocks to *policy functions*, that map the idiosyncratic state variables to equilibrium policies. To do so, we build on, and extend the operator learning approach of Zhong (2023).

5.3.1 Operator learning

Suppose that we want to obtain an approximation $\hat{f}(X, x)$ to the function $f(X, x)$, where $X \in \mathbb{R}^{N_1}$, $x \in \mathbb{R}^{N_2}$, and $f(X, x) \in \mathbb{R}$. Let $F(\mathbb{R}^{N_2}, \mathbb{R})$ denote the space of functions mapping $\mathbb{R}^{N_2}$ into $\mathbb{R}$. The idea of operator learning is to decompose the mapping f as $f(X, x) = g(X)(x)$. Here, the *operator* $g : X \to g(X) : \mathbb{R}^{N_1} \to F(\mathbb{R}^{N_2}, \mathbb{R})$, maps the vector $X \in \mathbb{R}^{N_1}$ to a *function* $g(X) \in F(\mathbb{R}^{N_2}, \mathbb{R})$. The function $g(X)$ is then evaluated at x.

For example, let X denote the aggregate state vector, let x denote the idiosyncratic state vector, and let $f(X, x)$ denote a policy function which depends on aggregate and idiosyncratic state variables. In this example $g : \mathbb{R}^{N_1} \to F(\mathbb{R}^{N_2}, \mathbb{R})$ maps the aggregate state to a policy function. The policy function $g(X)$ in turn maps the idiosyncratic state x to the choice variable $g(X)(x)$. Consider the case of $N_1 = N_2 = 1$ and suppose that, conditional on the aggregate state X, the policy function is a linear function of the idiosyncratic state x, that is, $f(X, x) = g(X)(x) = \alpha_X x$. Although the policy function is linear in the idiosyncratic state x, the slope coefficient α_X depends on the aggregate state X in a potentially nonlinear way. We can split the problem of approximating $f(X, x)$ into two parts. First, approximating $g(X)$ by $\hat{g}(X) : X \to \hat{\alpha}_X$.[37] Second, evaluating $\hat{g}(X)(x) = \hat{\alpha}_X x$. As we show in this paper, a key advantage of approximating operators, as opposed to directly parameterizing policy functions, is that it is easy to ensure desirable properties of $\hat{g}(X)(x)$, such as monotonicity or concavity in x for all X.

Shape-preserving operators For concreteness, let us consider the household consumption function from the model above: $\pi^c(\mathbf{X}_t^{\text{agg. state}}, e_t^i, \theta_t^i) = c_t^i$. Following the operator learning approach, we decompose the problem of predicting the individual policy into two steps: First we predict a *function* $\pi_t^{c,\text{ ind}}$ for each aggregate state $\mathbf{X}_t^{\text{agg. state}}$. Then we evaluate the resulting function $\pi_t^{c,\text{ ind}}(e_t^i, \theta_t^i) = c_t^i$ for all idiosyncratic states of interest.

Often, some crucial shape properties of the equilibrium functions $\pi_t^{c,\text{ ind}}(e_t^i, \theta_t^i)$ are known ex-ante. In the context of the model we present in this section, or for a large class of other macroeconomic models, we know that the consumption function is strictly increasing and concave in the idiosyncratic

[37] A linear function from $\mathbb{R} \to \mathbb{R}$ is fully characterized by its slope coefficient. Therefore, we can essentially predict a function by predicting a single real number.

asset holding θ^i.[38] In this section, we show a simple way to ensure that all predicted policy functions fulfill these properties. Thus, we effectively restrict the search space to the space of economically more meaningful functions, increasing the robustness of our algorithm. What is more, a monotone consumption function allows us to apply the endogenous grid method. Using the endogenous grid method, we obtain targets to train the policy function with supervised learning, increasing the robustness of our approach.

The first step is to choose a functional form to approximate the function $\pi_t^{\text{c, ind}}(e_t^i, \theta_t^i)$. We consider the case where e_t^i takes two discrete values $\mathcal{E}^{\text{grid}} = [e_0, e_1]$, where $e_0 < e_1$, and where θ_t^i is continuous. Because e_t^i is a discrete variable, we can split $\pi^{\text{c, ind}}$ into two univariate functions. Then, we approximate those functions using piecewise linear interpolation on a grid of interpolation nodes $\Theta^{\text{grid}} = [\theta_0, \theta_1, \dots, \theta_{N-1}]$, where $\theta_0 < \theta_1 < \dots < \theta_{N-1}$

$$\mathcal{C}_t^{\text{grid}} = \begin{bmatrix} c_{t,0,0}, c_{t,0,1}, \dots, c_{t,0,N-1} \\ c_{t,1,0}, c_{t,1,1}, \dots, c_{t,1,N-1} \end{bmatrix} \in \mathbb{R}^{2\times N}, \tag{76}$$

where $c_{t,i,j} = \pi_t^{\text{c, ind}}(e_i, \theta_j)$ for a given aggregate state $\mathbf{X}_t^{\text{agg. state}}$, idiosyncratic states $e_i \in \mathcal{E}^{\text{grid}}$, and $\theta_j \in \Theta^{\text{grid}}$. For asset holdings θ between grid points, *i.e.* $\theta_j < \theta < \theta_{j+1}$, we interpolate linearly between $c_{t,i,j}$ and $c_{t,i,j+1}$. Conditional on an aggregate state and fixed grids $\mathcal{E}^{\text{grid}}$ and Θ^{grid}, the approximate household consumption function is fully described by the $2 \times N$ values of $\mathcal{C}_t^{\text{grid}}$, which in turn are parameterized by a neural network. Following our sequence space approach, this neural network takes the truncated sequence of aggregate shocks $\mathbf{X}_t^{\text{seq},T}$ as an input.

Monotonicity: To ensure that the predicted consumption functions are increasing in idiosyncratic asset holdings θ, we need to make sure that the prediction by the neural network always satisfies $c_{t,i,j} < c_{t,i,j+1}\ \forall t, i, j$. To ensure that this condition holds, we follow a two-step procedure. First, instead of predicting $\mathcal{C}_t^{\text{grid}}$ directly, the neural network predicts the boundary consumption value[39] and consumption increments

$$\mathcal{N}_{\boldsymbol{\rho}}^c(\mathbf{X}_t^{\text{seq},T}) \approx d\mathcal{C}_t^{\text{grid}} := \begin{bmatrix} c_{t,0,0}, dc_{t,0,1}, \dots, dc_{t,0,N-1} \\ c_{t,1,0}, dc_{t,1,1}, \dots, dc_{t,1,N-1} \end{bmatrix} \in \mathbb{R}_{>0}^{2\times N}, \tag{77}$$

where we ensure that all the $2 \times N$ predicted entries of $d\mathcal{C}_t^{\text{grid}}$ are positive. This is easy to ensure, for example, by using a *softplus* activation function in the output layer. In the second step we construct

$$\mathcal{C}_t^{\text{grid}} = \begin{bmatrix} c_{t,0,0}, c_{t,0,1}, \dots, c_{t,0,N-1} \\ c_{t,1,0}, c_{t,1,1}, \dots, c_{t,1,N-1} \end{bmatrix}, \tag{78}$$

where

$$c_{t,i,j} = c_{t,i,0} + \sum_{k=1}^{j} dc_{t,i,k}. \tag{79}$$

[38] For example, this is the case in standard incomplete-market models in the spirit of Imrohoroğlu (1989); Bewley (1977); Huggett (1993); Aiyagari (1994); Krusell and Smith (1998).

[39] Consumption at the lowest gridpoint in the asset grid.

Since the architecture of our neural network guarantees that all the predicted increments $dc_{t,i,k}$ are positive, this two-step procedure guarantees monotonicity of the predicted consumption at the interpolation nodes: $c_{t,i,j} < c_{t,i,j+1}$. Piecewise linear interpolation guarantees that monotonicity of the grid values translates into monotonicity of the resulting function. A large number of gridpoints, N, and, for example, a log-log spaced grid[40] for the asset holdings Θ^{grid} can ensure that the loss in accuracy coming from the linear interpolation is minimal.

Concavity: In order to jointly guarantee monotonicity and concavity of the consumption function, one needs to ensure that the slope of the consumption function is decreasing in the individual asset holding, while it remains bounded from below by zero. Let $d\theta_j := \theta_j - \theta_{j-1}$ define the distance between two grid points θ_j and θ_{j-1} on the asset grid Θ^{grid}. To ensure concavity and monotonicity of consumption, we follow a three-step procedure. In a first step, we predict:

$$\mathcal{N}^c_{\boldsymbol{\rho}}(\mathbf{X}_t^{\text{seq},T}) \approx dd\mathcal{C}_t^{\text{grid}} := \begin{bmatrix} c_{t,0,0}, ddc_{t,0,1}, ddc_{t,0,2} \ldots, ddc_{t,0,N-1} \\ c_{t,1,0}, ddc_{t,1,1}, ddc_{t,1,2} \ldots, ddc_{t,1,N-1} \end{bmatrix}, \tag{80}$$

where we ensure that $c_{t,i,0} > 0$ and $ddc_{t,i,j} < 0$ using a suitable activation function. In the second step, we construct the matrix $d\mathcal{C}_t^{\text{grid}}$, using

$$dc_{t,i,j} = d\theta_j \times exp\left(\sum_{k=1}^{j} ddc_{t,i,k}\right). \tag{81}$$

Since all predictions $ddc_{t,i,k} < 0$, our architecture enforces that $dc_{t,i,j}/d\theta_j$ is decreasing in j, ensuring concavity of the consumption function. The exponential function further guarantees that all $dc_{t,i,j}/d\theta_j > 0$, making sure that the consumption function is also monotone. In the third step, we use the cumulative sum to construct the implied consumption grid, as in equation (79).

Borrowing constraints: With a slight modification of the procedure described above, we can additionally encode the borrowing constraint $\theta^i_{t+1} \geq \underline{\theta} = 0$ directly into the network architecture. For a given asset price and given dividends, we can convert the distance between points in the asset grid, $d\theta_j$, to differences in cash-at-hand (cah) by multiplying the asset grid distance with the payout of the asset.[41] In our model: $dcah_j = d\theta_j(D_t + \Gamma p_t)$. We know that the consumption function is increasing and concave in cash-at-hand. We can ensure monotonicity and concavity of the consumption function by making sure that the marginal propensity to consume, *i.e.* $dc/dcah$, is positive and decreasing. To additionally ensure that the borrowing constraint is always satisfied, it is sufficient to ensure that the consumption share at the borrowing constraint and the marginal propensities to consume are bounded from above by 1. We can achieve this by following a three-step procedure, similar to the one outlined in the previous paragraph. The neural network now predicts

$$\mathcal{N}^c_{\boldsymbol{\rho}}(\mathbf{X}_t^{\text{seq},T}) \approx d\mathcal{MPC}_t^{\text{grid}} := \begin{bmatrix} cs_{t,0,0}, dmpc_{t,0,1}, dmpc_{t,0,2} \ldots, dmpc_{t,0,N-1} \\ cs_{t,1,0}, dmpc_{t,1,1}, dmpc_{t,1,2} \ldots, dmpc_{t,1,N-1} \end{bmatrix}, \tag{82}$$

[40] *Log-log* ensures high grid density for low asset holdings, where the short sale limit might bind.

[41] The income form other sources, such as labor income and startup rents, is constant across the asset grid, and can hence be omitted in this calculation.

where we use a sigmoid activation function to ensure that $0 < cs_{t,i,0} \leq 1$. Where $cs_{t,i,0}$ denotes the consumption share out of cash-at-hand at the first grid point on the asset grid. We use a softplus activation and multiplication with -1 to ensure that all predicted $dmpc_{t,i,j} < 0$. In the second step we construct the consumption value at the first asset grid point $c_{t,i,0} = cs_{t,i,0} \times cah_{t,i,0}$. Since $cs_{t,i,0}$ always lies between 0 and 1, the borrowing constraint is satisfied. We then obtain the marginal propensity to consume

$$mpc_{t,i,j} = \underbrace{cs_{t,i,0}}_{\in(0,1)} \times \underbrace{e^{\sum_{k=1}^{j} dmpc_{t,i,k}}}_{\in(0,1),\text{ decr in j}}, \tag{83}$$

for $j > 0$ and obtain

$$\mathcal{MPC}_t^{\text{grid}} := \begin{bmatrix} c_{t,0,0}, mpc_{t,0,1}, mpc_{t,0,2}, \ldots, mpc_{t,0,N-1} \\ c_{t,1,0}, mpc_{t,1,1}, mpc_{t,1,2}, \ldots, mpc_{t,1,N-1} \end{bmatrix}. \tag{84}$$

This parameterization ensures that all predicted $mpc_{t,i,j}$ are positive, bounded from above by 1, and decreasing in the index j. In the third step, we again compute the remaining consumption values for all gridpoints

$$c_{t,i,j} = c_{t,i,0} + \sum_{k=1}^{j} mpc_{t,i,k} \times dcah_{t,i,k}. \tag{85}$$

As in the case of parameterizing *ddc*, the resulting consumption grid is increasing and concave. Because the predicted marginal propensities to consume are bounded from above by 1, this construction additionally ensures that the borrowing constraint is never violated.[42] It also makes it easy for the neural network to predict the consumption for borrowing constraint households extremely precisely (by predicting the maximum mpc of 1). In models, where the marginal propensity to consume plays a central role, ensuring a precisely satisfied borrowing constraint can be particularly important for the implied model predictions.

Shape-preservation in other settings: It is worth noting that our approach of shape-preserving operator learning is not limited to the sequence-space approach or to discrete-time models. Shape-preserving operator learning can be analogously applied in state-space-based methods or in continuous-time models. Lastly, note that our architecture guarantees only the monotonicity and concavity of the predicted values at the interpolation nodes. With linear interpolation in one dimension, the monotonicity and concavity of the values at the interpolation nodes are sufficient for the monotonicity and (weak) concavity of the interpolating function. Shape-preserving interpolation of multivariate functions is a more complicated problem. Multivariate piecewise linear interpolation is guaranteed to preserve monotonicity; however, concavity of the interpolating function is generally not guaranteed for simple interpolation methods in higher dimensions.

[42] Given that we ensure that the consumption value at the first gridpoint is consistent with the borrowing constraint.

5.3.2 Firm policies

As part of computing an approximate recursive equilibrium, we need to solve for the firms' policy functions $\pi^F(\mathbf{X}_t^{\text{agg. state}}, z_t^i, k_t^i) = k_{t+1}^i$ and $\pi^\lambda(\mathbf{X}_t^{\text{agg. state}}, z_t^i, k_t^i) = \lambda_t^i$, such that these functions are consistent with the optimality conditions given in equations (51) to (54). We approximate both functions using the operator learning approach. Employing the shape-preservation techniques described above, we make sure that the predicted KKT multiplier is non-negative and weakly decreasing in firm capital k. Similarly, we ensure that the firm savings function is increasing in k.

The firm-level productivity follows a three-state Markov chain. We set a grid of N_k points over the firm capital holdings. For each aggregate shock, we obtain the following predictions

$$\mathcal{K}_t^{\text{grid}} = \begin{bmatrix} k_{t+1,0,0} & k_{t+1,0,1} & \dots & k_{t+1,0,N_k-1} \\ k_{t+1,1,0} & k_{t+1,1,1} & \dots & k_{t+1,1,N_k-1} \\ k_{t+1,2,0} & k_{t+1,2,1} & \dots & k_{t+1,2,N_k-1} \end{bmatrix} \tag{86}$$

$$\lambda_t^{\text{grid}} = \begin{bmatrix} \lambda_{t,0,0} & \lambda_{t,0,1} & \dots & \lambda_{t,0,N_k-1} \\ \lambda_{t,1,0} & \lambda_{t,1,1} & \dots & \lambda_{t,1,N_k-1} \\ \lambda_{t,2,0} & \lambda_{t,2,1} & \dots & \lambda_{t,2,N_k-1} \end{bmatrix}, \tag{87}$$

corresponding to the gridded idiosyncratic firm policies for a given aggregate state at the $3 \times N_k$ interpolation nodes. To evaluate the policies for capital values between grid points, we rely on linear interpolation.

We use a shape-preserving operator network to ensure that the predicted values for next period's capital, $k_{t+1,i,j}$ are increasing in j (*i.e.* increasing in $k_{t,i,j}$). For the KKT multiplier, $\lambda_{t,i,j}$, we similarly ensure that all the predictions are non-negative and weakly decreasing in j. Following our sequence space approach, the input to the neural networks is given by a truncated sequence of aggregate shocks

$$\mathbf{X}_t^{\text{seq},T} = [\underbrace{\epsilon_{t-T+1}^A, \epsilon_{t-T+2}^A, \dots, \epsilon_t^A}_{\text{shocks to aggregate TFP}}, \underbrace{\chi_{t-T+1}, \chi_{t-T+2}, \dots, \chi_t}_{\text{uncertainty regime}}]. \tag{88}$$

In the first step, the neural networks predict

$$\mathcal{N}_{\boldsymbol{\rho}}^k(\mathbf{X}_t^{\text{seq},T}) = \begin{bmatrix} k_{t+1,0,0} & dk_{t+1,0,1} & dk_{t+1,0,2} & \dots & dk_{t+1,0,N_k-1} \\ k_{t+1,1,0} & dk_{t+1,1,1} & dk_{t+1,1,2} & \dots & dk_{t+1,1,N_k-1} \\ k_{t+1,2,0} & dk_{t+1,2,1} & dk_{t+1,2,2} & \dots & dk_{t+1,2,N_k-1} \end{bmatrix} \tag{89}$$

$$\mathcal{N}_{\boldsymbol{\rho}}^\lambda(\mathbf{X}_t^{\text{seq},T}) = \begin{bmatrix} \log(\lambda_{t,0,0}) & d\lambda_{t,0,1} & d\lambda_{t,0,2} & \dots & d\lambda_{t,0,N_k-1} \\ \log(\lambda_{t,1,0}) & d\lambda_{t,1,1} & d\lambda_{t,1,2} & \dots & d\lambda_{t,1,N_k-1} \\ \log(\lambda_{t,2,0}) & d\lambda_{t,2,1} & d\lambda_{t,2,2} & \dots & d\lambda_{t,2,N_k-1}, \end{bmatrix} \tag{90}$$

where we use the softplus activation in the output layer to ensure that $k_{t+1,i,j} \geq 0$, $dk_{t+1,i,j} \geq 0$, and $-$softplus to ensure that $d\lambda_{t,i,j} < 0$. In the second step, we construct $\mathcal{K}_{t+1}^{\text{grid}}$ and λ_t^{grid} to ensure

non-negativity and monotonicity of both functions:

$$k_{t+1,i,j} = k_{t+1,i,0} + \sum_{h=1}^{j} dk_{t+1,i,h} \tag{91}$$

$$\lambda_{t,i,j} = exp\left(\log(\lambda_{t,i,0}) + \sum_{h=1}^{j} d\lambda_{t,i,h}\right). \tag{92}$$

Simulation We represent the cross-sectional distribution of firms over idiosyncratic capital and productivity using the finite histogram method of Young (2010). We base the firm histogram on the same grid as we use for interpolation nodes of firms' policy functions. Because of that, we can evaluate the histogram transition using the gridded capital policies predicted by the neural network without the need to resort to interpolation.

Loss function To evaluate the performance of our neural network in producing policies consistent with firm optimality conditions, we need to define a loss function to map equilibrium conditions error into a scalar quantity. As in Azinovic et al. (2022), we choose the standard mean-squared function as our loss. Through the stochastic discount factor the firms' optimality conditions also depend on household policies and on the equity price. Because of that, the loss function also depends on the household network $\mathcal{N}^c_{\boldsymbol{\rho}}$, and on the price network $\mathcal{N}^p_{\boldsymbol{\rho}}$.

We use Gauss-Hermite quadrature with N^{GH} quadrature nodes to approximate the expectation over realizations of shocks to aggregate productivity next period, for each of the two uncertainty regimes. Let $\mathcal{A}_t := \left(\mathbf{X}_t^{\text{state}}, \{\mathbf{X}_{t+1,i}^{\text{state}}\}_{i=1}^{2N_{GH}}, \mathbf{X}_t^{\text{seq},T}, \mathbf{X}_{t+1,i}^{\text{seq},T}\}_{i=1}^{2N_{GH}}\right)$, denote aggregate states and truncated histories of aggregate shocks for a given realization in period t, and at all $2N_{GH}$ integration nodes for the next period $t+1$. Given $\mathcal{A}_t$, the wage can be computed as a closed form function of the state-vector, as given in equation (58). The outputs of the firm network $\mathcal{N}^k_{\boldsymbol{\rho}}$, the outputs of the household network $\mathcal{N}^c_{\boldsymbol{\rho}}$, and the output of the price network $\mathcal{N}^p_{\boldsymbol{\rho}}$ allows us to construct the aggregate stochastic discount factor $\Lambda_{t+1,t}$, aggregate dividends $\left(D_t, \{D_{t+1,i}\}_{i=1}^{2N_{GH}}\right)$, as well as the stock prices $\left(p_t, \{p_{t+1,i}\}_{i=1}^{2N_{GH}}\right)$. We then sample idiosyncratic states (z_t^i, k_t^i),[43] and evaluate the firm policies on those states to obtain $k_{t+1}^i, \{k_{t+2,j}^i\}_{j=1}^{2N_{GH}}$, and $\{\lambda_{t+1,j}^i\}_{j=1}^{2N_{GH}}$. Then we recover the implied KKT multiplier, $\lambda_t^{i,\text{ implied}}$, from the optimality condition (51):

$$\lambda_t^{i,\text{ implied}} = \frac{\Gamma \mathrm{E}\left[\Lambda_{t+1,t}\left(1 + r_{t+1}^{i,k} + \frac{\partial}{\partial k_{t+1}^i}\psi(k_{t+2}^i, k_{t+1}^i)\right)(1+\lambda_{t+1}^i)\right]}{\left(1 + \frac{\partial}{\partial k_{t+1}^i}\psi(k_{t+1}^i, k_t^i)\right)} - 1. \tag{93}$$

The difference between the predicted and the implied Lagrange multiplier summarizes the error in the firm Euler equation. Besides the Euler equation error, we also need to evaluate the error in the

[43] We sample the idiosyncratic labor endowment states from the ergodic distribution of the corresponding Markov chain. For the idiosyncratic capital, we sample half of the points from the capital grid, and the other half from a uniform distribution.

remaining KKT conditions.

$$err_1^{\text{firm}}\left(\mathcal{A}_t, \mathcal{N}_{\boldsymbol{\rho}}^c, \mathcal{N}_{\boldsymbol{\rho}}^p, \mathcal{N}_{\boldsymbol{\rho}}^k, \mathcal{N}_{\boldsymbol{\rho}}^\lambda, z_t^i, k_t^i\right) = \frac{\lambda_t^i - \lambda_t^{i,\text{ implied}}}{1 + \lambda_t^{i,\text{ implied}}} \tag{94}$$

$$err_2^{\text{firm}}\left(\mathcal{A}_t, \mathcal{N}_{\boldsymbol{\rho}}^c, \mathcal{N}_{\boldsymbol{\rho}}^p, \mathcal{N}_{\boldsymbol{\rho}}^k, \mathcal{N}_{\boldsymbol{\rho}}^\lambda, z_t^i, k_t^i\right) = \psi^{FB}\left(\frac{d_t^i - \underline{d}}{k_t^i}, \lambda_t^{i,\text{ implied}}\right), \text{ where} \tag{95}$$

$$\psi^{FB}(a,b) = a - b - \sqrt{a^2 + b^2}, \tag{96}$$

denotes the Fischer-Burmeister function (see, *e.g.* Jiang, 1999). The optimal firm policies are characterized by

$$err_1^{\text{firm}}\left(\mathcal{A}_t, \mathcal{N}_{\boldsymbol{\rho}}^c, \mathcal{N}_{\boldsymbol{\rho}}^p, \mathcal{N}_{\boldsymbol{\rho}}^k, \mathcal{N}_{\boldsymbol{\rho}}^\lambda, z_t^i, k_t^i\right) = 0 \tag{97}$$

$$err_2^{\text{firm}}\left(\mathcal{A}_t, \mathcal{N}_{\boldsymbol{\rho}}^c, \mathcal{N}_{\boldsymbol{\rho}}^p, \mathcal{N}_{\boldsymbol{\rho}}^k, \mathcal{N}_{\boldsymbol{\rho}}^\lambda, z_t^i, k_t^i\right) = 0, \tag{98}$$

for all $\mathcal{A}_t$ and all z_t^i, k_t^i.

5.3.3 Household policies and stock price

The two additional equilibrium functions, which we need to compute are the household policy function, and the equity price function. For households, we choose to parameterize the consumption function $\pi^C(\mathbf{X}_t^{\text{agg. state}}, e_t^i, \theta_t^i) = c_t^i$. For the stock price we need to approximate $\pi^P(\mathbf{X}^{\text{agg. state}}) = p_t$. We parameterize the equity price using a simple densely connected feed-forward neural network that maps the truncated history of aggregate shocks into a non-negative equity price[44]

$$\mathcal{N}_{\boldsymbol{\rho}}^p(\mathbf{X}_t^{\text{seq},T}) = p_t. \tag{99}$$

To parameterize household consumption function, we use our shape-preserving neural network operator architecture, to obtain a matrix $\mathcal{C}_t^{\text{grid}}$, that represents the consumption function on a discrete grid over the idiosyncratic assets θ_j and the labor endowment e_i:

$$\mathcal{C}_t^{\text{grid}} = \begin{bmatrix} c_{t,0,0}, c_{t,0,1}, \ldots, c_{t,0,N} \\ c_{t,1,0}, c_{t,1,1}, \ldots, c_{t,1,N} \end{bmatrix}, \tag{100}$$

where $c_{t,i,j}$ denote the consumption in period t for each idiosyncratic productivity level e_i in the productivity grid $\mathcal{E}^{\text{grid}} = [e_0, e_1]$, and for each asset holding θ_j in the asset grid $\Theta^{\text{grid}} = [\theta_0, \theta_1, \ldots, \theta_{N_\theta - 1}]$, with $\theta_0 < \theta_1 < \cdots < \theta_{N_\theta - 1}$. We use our shape-preserving neural network architecture to ensure that the resulting consumption function satisfies three properties: first, it is increasing in cash-at-hand, second, it is concave in cash-at-hand, and third, it is consistent with the borrowing constraint $\theta_{t+1} \geq \underline{\theta} = 0$. We impose these three restrictions by following the three-step

[44] To make sure the network predicts non-negative price, we activate the output layer with a *softplus* activation function.

procedure, outlined in section 5.3.1. The neural network learns to predict

$$\mathcal{N}^c_{\boldsymbol{\rho}}(\mathbf{X}_t^{\text{seq},T}) \approx d\mathcal{MPC}_t^{\text{grid}} := \begin{bmatrix} cs_{t,0,0}, dmpc_{t,0,1}, dmpc_{t,0,2} \ldots, dmpc_{t,0,N} \\ cs_{t,1,0}, dmpc_{t,1,1}, dmpc_{t,1,2} \ldots, dmpc_{t,1,N} \end{bmatrix}, \tag{101}$$

where $cs_{t,i,0}$ denotes the consumption share out of cash-at-hand at the first grid point on the asset grid. As outlined in section 5.3.1, the neural network architecture ensures that all predicted $dmpc_{t,i,j} < 0$, as well as $0 < cs_{t,i,0} \leq 1$. In a second step we then construct

$$c_{t,i,0} = cs_{t,i,0} \times cah_{t,i,0} \tag{102}$$

$$mpc_{t,i,j} = cs_{t,i,0} \times e^{\sum_{k=1}^{j} dmpc_{t,i,k}}, \tag{103}$$

for $j > 0$ and obtain $\mathcal{MPC}_t^{\text{grid}}$. By construction, all predicted $mpc_{t,i,j}$ are positive, bounded from above by 1, and decrease in the index j. In the third step, we construct the consumption values for all the remaining grid points

$$c_{t,i,j} = c_{t,i,0} + \sum_{k=1}^{j} mpc_{t,i,k} \times dcah_{t,i,k} \tag{104}$$

$$dcah_{t,i,j} = d\theta_{t,i,j}(D_t + \Gamma p_t) \tag{105}$$

completing the construction of $\mathcal{C}_t^{\text{grid}}$.

Loss function Instead of backpropagating through the complex computational graph defined by the forward-looking optimality condition, we use the method of endogenous gridpoints (see Carroll, 2006),[45] to derive the period t consumption function implied by the current guess of equilibrium functions in period $t+1$. Exploiting speed and robustness of the endogenous gridpoint method, we are furthermore able to use a simple Newton-Raphson method to solve for market clearing equity price p_t^{solved} together with implied consumption function $\mathcal{C}_t^{\text{grid, solved}}$. Using this procedure, we obtain *training targets* for price and household networks, allowing us to train these networks on a simple supervised learning loss.

From $\mathcal{C}_t^{\text{grid, solved}}$ and p_t^{solved}, we can also obtain target values for the intermediate predictions, *i.e.* $\mathcal{MPC}_t^{\text{grid, solved}}$. The supervised learning loss terms for household and price blocks are summarized in the following equations:

$$\text{err}_1^{\text{hh}}\left(\mathcal{A}_t, \mathcal{N}^c_{\boldsymbol{\rho}}, \mathcal{N}^p_{\boldsymbol{\rho}}, \mathcal{N}^k_{\boldsymbol{\rho}}\right) = \left(\mathcal{C}_t^{\text{grid}} - \mathcal{C}_t^{\text{grid, solved}}\right) \oslash \mathcal{C}_t^{\text{grid, solved}} \tag{106}$$

$$\text{err}_2^{\text{hh}}\left(\mathcal{A}_t, \mathcal{N}^c_{\boldsymbol{\rho}}, \mathcal{N}^p_{\boldsymbol{\rho}}, \mathcal{N}^k_{\boldsymbol{\rho}}\right) = \left(\mathcal{MPC}_t^{\text{grid}} - \mathcal{MPC}_t^{\text{grid, solved}}\right) \oslash \mathcal{MPC}_t^{\text{grid, solved}} \tag{107}$$

$$\text{err}_1^{\text{p}}\left(\mathcal{A}_t, \mathcal{N}^c_{\boldsymbol{\rho}}, \mathcal{N}^p_{\boldsymbol{\rho}}, \mathcal{N}^k_{\boldsymbol{\rho}}\right) = \frac{p_t - p_t^{\text{solved}}}{p_t^{\text{solved}}}, \tag{108}$$

where $\oslash$ denotes element-wise division, and p_t, $\mathcal{C}_t^{\text{grid}}$, and $\mathcal{MPC}_t^{\text{grid}}$ denote predictions by the neural

[45] For the endogenous grid method in higher dimensions, see Druedahl and Jørgensen (2017).

networks.

5.3.4 Loss function

Our overall loss function to train the neural networks is given by

$$\begin{aligned}\mathcal{L}(\boldsymbol{\rho}, \mathcal{A}_t) = & w_1^{\text{firm}} \left(\frac{1}{N_{\text{sample}}} \sum_{i=1}^{N_{\text{sample}}} \text{err}_1^{\text{firm}} \left(\mathcal{A}_t, \mathcal{N}_{\boldsymbol{\rho}}^c, \mathcal{N}_{\boldsymbol{\rho}}^p, \mathcal{N}_{\boldsymbol{\rho}}^k, \mathcal{N}_{\boldsymbol{\rho}}^\lambda, z_t^i, k_t^i \right)^2 \right) \\ & + w_2^{\text{firm}} \left(\frac{1}{N_{\text{sample}}} \sum_{i=1}^{N_{\text{sample}}} \text{err}_2^{\text{firm}} \left(\mathcal{A}_t, \mathcal{N}_{\boldsymbol{\rho}}^c, \mathcal{N}_{\boldsymbol{\rho}}^p, \mathcal{N}_{\boldsymbol{\rho}}^k, \mathcal{N}_{\boldsymbol{\rho}}^\lambda, z_t^i, k_t^i \right)^2 \right) \\ & + w_1^{\text{hh}} \left(\frac{1}{2N_\theta} \sum_{i=0}^{1} \sum_{j=0}^{N_\theta - 1} \left[\text{err}_1^{\text{hh}} \left(\mathcal{A}_t, \mathcal{N}_{\boldsymbol{\rho}}^c, \mathcal{N}_{\boldsymbol{\rho}}^p, \mathcal{N}_{\boldsymbol{\rho}}^k \right) \right]_{i,j}^2 \right) \\ & + w_2^{\text{hh}} \left(\frac{1}{2N_\theta} \sum_{i=0}^{1} \sum_{j=0}^{N_\theta - 1} \left[\text{err}_2^{\text{hh}} \left(\mathcal{A}_t, \mathcal{N}_{\boldsymbol{\rho}}^c, \mathcal{N}_{\boldsymbol{\rho}}^p, \mathcal{N}_{\boldsymbol{\rho}}^k \right) \right]_{i,j}^2 \right) \\ & + w^p \left(\text{err}_1^{\text{p}} \left(\mathcal{A}_t, \mathcal{N}_{\boldsymbol{\rho}}^c, \mathcal{N}_{\boldsymbol{\rho}}^p, \mathcal{N}_{\boldsymbol{\rho}}^k \right) \right)^2, \end{aligned} \tag{109}$$

where w_1^{firm}, w_2^{firm}, w_1^{hh}, w_2^{hh}, and w^p denote the weights of the different components of the loss function and where, with slight abuse of notation, $\boldsymbol{\rho}$ collects the trainable parameters of all four neural networks.

5.3.5 Forward simulation

Given the aggregate state-history pair, $\mathcal{A}_t$, we obtain $\mathcal{A}_{t+1}$ by drawing new random shocks for productivity and for the uncertainty regime according to the corresponding probability distributions, and by evolving the endogenous state according to the current guess of equilibrium policy functions. The evolution of the distribution of firms is computed using policy functions obtained by evaluation of the firm network $\mathcal{N}_{\boldsymbol{\rho}}^k$. To make sure that the evolution of the distribution of households satisfies market clearing condition, we use the market clearing policies $\mathcal{C}_t^{\text{grid, solved}}$ and p_t^{solved} computed using a Newton-Raphson root-finding algorithm.[46]

Computing the target values In order to compute the target values for household consumption, and for the stock price, $\mathcal{C}_t^{\text{grid, solved}}$ and p_t^{solved}, that are consistent with household optimality and market clearing, we rely on the endogenous gridpoint method (EGM). Specifically, given a guess for current stock price p_t^{guess}, and current guess for equilibrium dynamics of the economy,[47] we use the EGM to derive the period t consumption function implied by expectations of $t+1$ marginal utility of consumption and asset payouts. We denote this consumption function $c_t^{i,\text{ guess}}$ and corresponding

[46] Note that the consumption policies, together with the stock price, pin down the asset policies via the households' budget constraint.

[47] *i.e.* we evaluate the evolution of the firm distribution using the firm policy functions encoded in the firm network, and then we evaluate the evolution of the household distribution using household network as an input to solve for market clearing price and policy.

savings function as $\theta_{t+1}^{i,\,\text{guess}}$. Next, we evaluate the excess demand function implied by household policies generated by the current price guess $ED(p_t^{\text{guess}}, \mathcal{A}_t, \mathcal{N}_{\boldsymbol{\rho}}^c, \mathcal{N}_{\boldsymbol{\rho}}^p, \mathcal{N}_{\boldsymbol{\rho}}^k) = \Theta_{t+1}^{\text{guess}} - 1$. Then we update the price guess using a Newton-Raphson algorithm

$$p_t^{\text{new}} = p_t^{\text{old}} - \frac{ED(p_t^{\text{old}}, \mathcal{A}_t, \mathcal{N}_{\boldsymbol{\rho}}^c, \mathcal{N}_{\boldsymbol{\rho}}^p, \mathcal{N}_{\boldsymbol{\rho}}^k)}{\frac{\partial}{\partial p_t^{\text{old}}} ED(p_t^{\text{old}}, \mathcal{A}_t, \mathcal{N}_{\boldsymbol{\rho}}^c, \mathcal{N}_{\boldsymbol{\rho}}^p, \mathcal{N}_{\boldsymbol{\rho}}^k)}. \tag{110}$$

Our shape-preserving neural network architecture ensures that the excess demand function is well behaved in the price guess. Therefore, a small number of Newton-Raphson steps is sufficient to compute the market clearing price.[48] Furthermore, the prediction of the neural network $\mathcal{N}_{\boldsymbol{\rho}}^p$ provides us with an excellent starting guess for the market clearing price, once the initial training stage is completed.[49] When the market clearing price p_t^{solved} is found, we save the price and the corresponding household policies $\mathcal{C}_t^{\text{grid, solved}}$ as target for supervised training.

5.4 Training

5.4.1 Outline of the step-wise training procedure

We again follow a step-wise training procedure. Our step-wise procedure ensures that our algorithm learns reasonable firm policies, and hence generates sensible wages and dividends before it starts solving the household block.

Step 1: Firm side only: In the first step we only train the firm side of the model. To do so, we set the stochastic discount factor of the firm to be equal to the patience parameter $\Lambda_{t+1,t} = \beta$, and set the weights in the loss function to $w_1^{\text{firm}} = w_2^{\text{firm}} = 1$ and $w_1^{\text{hh}} = w_2^{\text{hh}} = w^p = 0$. Imposing an exogenous stochastic discount factor and setting weights on household and price error to zero allows us to isolate and solve the firm problem, so we can then start solving the household problem with a good guess for firm policies already at hand. Furthermore, we start the training procedure with lower values for the adjustment cost parameters $\xi^{\text{up}} = 0.1$ and $\xi^{\text{down}} = 0.25$. The policies for the simplified parameterization can be learned quickly by the neural network.

Step 2, Pre-training price and household policies In the second step, we ensure that the neural networks parameterizing household policy and the equity price function also start with a good initial guess. To facilitate initial convergence, we again start the training procedure in a simplified economy: we introduce an artificial parameter τ. τ modifies the payout of equity to $p_{t+1}\Gamma(1-\tau)+D_{t+1}$. With $\tau = 0$, we recover the original economy. Setting $\tau = 1$ makes equity effectively a claim on $t+1$ dividends, reducing it into a short-lived asset. Given $\tau = 1$, we use $p_t^{\text{pre-train}} = D_t$ as a pre-training target for the price network $\mathcal{N}_{\boldsymbol{\rho}}^p$. For the household policy function, we use $\theta_{t+1}^{\text{i}} = \max(0, 0.7\theta_t^i - 0.1)$ for $e_t^i = e_0$ and $\theta_{t+1}^{\text{i}} = 0.7\theta_t^i + 0.6$ for $e_t^i = e_1$ as pre-training targets. This provides us with a good starting guess for the price function (in the short-lived asset calibration, $\tau = 1$), and the household policy functions.

[48] For GPU efficiency reasons, we use a predetermined number of Newton-Raphson steps, so that the same number of steps is used for all the states. The derivative of the excess demand function with respect to the guessed price can be computed efficiently using automatic differentiation.

[49] Hence, the algorithm might start with the initial number of Newton-Raphson iterations set to 10, which can be reduced to 3 once the algorithm converges to sufficiently low level of price prediction errors.

Step 3, Training the firm and household side together: Retaining the simplified parameterization of the model, we now train all price and policy functions on the full equilibrium loss function, given in equation (109) with weights set to $w_1^{\text{firm}} = w_2^{\text{firm}} = w_1^{\text{hh}} = w_2^{\text{hh}} = w^p = 1$.

Step 4, Step-wise model transformation to full parameterization: We then gradually change the parameters to the desired parameterization of the full model. We linearly increase the adjustment cost parameters to the final values of $\xi^{\text{up}} = 1.0$ and $\xi^{\text{down}} = 2.5$. At the same time, we decrease τ from $\tau = 1$ to $\tau = 0$ following a quadratic schedule. Additionally, we gradually change the stochastic discount factor used by the firms from β to $\Lambda_{t+1,t}$. At each step, firms use a weighted average between β and $\Lambda_{t+1,t}$ as a stochastic discount factor. We start with a full weight of 1 on β. Then we gradually decrease it to 0 while we increase the weight on $\Lambda_{t+1,t}$ from 0 to 1.

Step 5, Training with the final model parameters: We train the neural networks on the full loss function with the final model parameterization, until we reach a sufficiently low level of remaining errors in equilibrium conditions.

5.4.2 Hyperparameters

We summarize the hyperparameters that we used to solve the model in table A6. We choose all four neural networks, $\mathcal{N}_{\boldsymbol{\rho}}^k$, $\mathcal{N}_{\boldsymbol{\rho}}^\lambda$, $\mathcal{N}_{\boldsymbol{\rho}}^c$, and $\mathcal{N}_{\boldsymbol{\rho}}^p$, to be densely connected feed-forward neural networks with three hidden layers and gelu activation functions. We truncate the history of shocks after 300 periods. The input layers consist of 600 neurons, corresponding to the truncated history of uncertainty regimes and the innovations to productivity. Each hidden layer consists of 1024 gelu activated neurons.

5.5 Accuracy

In this section, we demonstrate the accuracy of our solution. Since a closed form or a high-quality numerical reference solution is not available, we investigate the errors in the equilibrium conditions, implied by the price and policy functions encoded in the learned parameters of our neural networks. We report the accuracy measures in economically interpretable units. In addition, we provide a comprehensive set of statistics on the distribution of errors across the state space.

5.5.1 Firm policies

The optimal firm policies are characterized by the set of firms' KKT conditions, given in equations (51) to (54). Rearranging equation (51) we obtain the following:

$$0 = \frac{\Gamma \mathrm{E}\left[\Lambda_{t+1,t}\left(1 + r_{t+1}^{i,k} + \frac{\partial}{\partial k_{t+1}^i}\psi(k_{t+2}^i, k_{t+1}^i)\right)\left(1 + \lambda_{t+1}^i\right)\right]}{\left(1 + \frac{\partial}{\partial k_{t+1}^i}\psi(k_{t+1}^i, k_t^i)\right)\left(1 + \lambda_t^i\right)} - 1, \tag{111}$$

where the numerator in the first term denotes the expected marginal benefit of additional capital in the next period, and the denominator gives the marginal cost in the current period. The errors in equation (111) are therefore interpretable as relative marginal cost errors.[50] The left panel of figure 4

[50] An error of 0.01, for example, would indicate that the marginal benefit is 1% higher than the marginal cost.

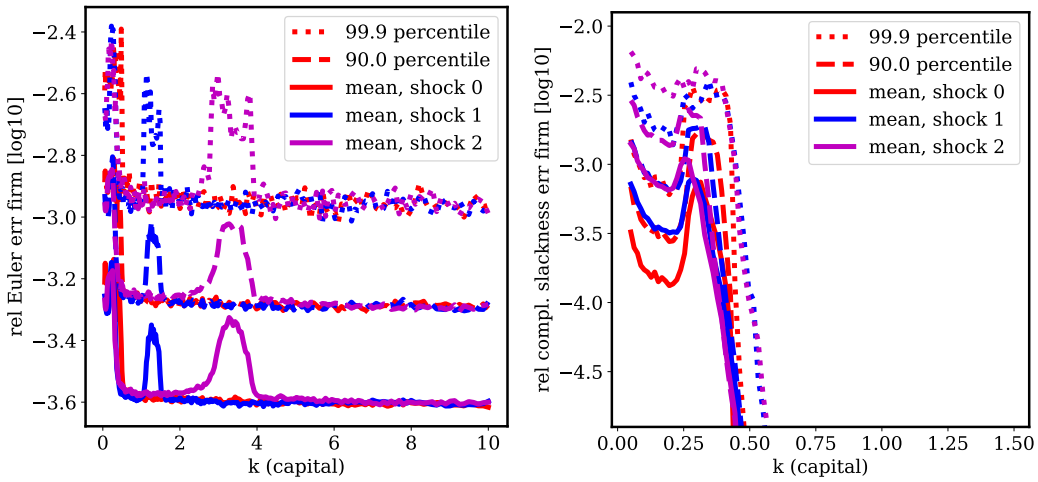


Figure 4: Remaining errors in the equilibrium conditions for the firm problem. The left panel shows the errors in the firms' Euler equations and the right panel shows the errors in the remaining KKT conditions.

shows statistics of the distribution of errors in equation (111) for different idiosyncratic productivity shocks, idiosyncratic levels of capital, and 1024 aggregate states drawn from the ergodic distribution of the economy. The mean marginal cost error remains below 0.07%, the 90th percentile below 0.16%, and the 99.9th percentile below 0.4% for all levels of firm capital and the three productivity levels. Even though the errors are very low everywhere, they are slightly larger in the areas of the state space where the firm policies feature a kink due to the constraint on dividends and due to asymmetric adjustment costs.[51]

The right panel of figure 4 shows statistics of the distribution of errors in the remaining KKT condition of firms

$$0 = \psi^{\text{FB}}\left(\frac{\lambda_t^i}{1+\lambda_t^i}, \frac{d_t^i - \underline{d}}{k_t^i}\right), \tag{112}$$

where $\psi^{\text{FB}}(x, y)$ denotes the Fischer-Burmeister function. The errors in equation (112) summarize the violations of all the remaining KKT conditions associated with the firm problem.[52] Except for very low levels of capital, the error is virtually zero. Throughout, the mean error remains below 0.2% and even the 99.9th percentile of errors remains below 0.8%. We conclude that firm policies are computed with high accuracy across the aggregate and idiosyncratic state space.

5.5.2 Household policies

As described above, the endogenous gridpoints method combined with a Newton-Raphson loop allows us to jointly solve for the market clearing equity price and the corresponding consumption

[51] For a fixed idiosyncratic level of capital and a fixed idiosyncratic productivity level, the mean and the percentiles are computed across 1024 aggregate states drawn from the ergodic distribution of the economy.

[52] Because the Fischer-Burmeister function satisfies $\psi^{\text{FB}}(x, y) = 0 \Leftrightarrow x \geq 0, y \geq 0, xy = 0$.

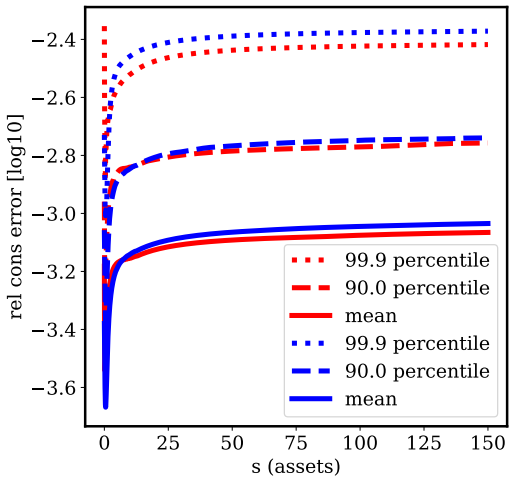


Figure 5: Remaining errors in the equilibrium conditions for the households' optimality conditions, expressed in units of relative consumption errors.

policies of households. Hence, we can assess the accuracy of the consumption policies learned by the neural network by comparing the neural network predictions for period t to the period t consumption policies implied by household optimality and market clearing prices.[53] Figure 5 shows the accuracy of the consumption function learned by the neural network across the aggregate and idiosyncratic state space. For all the idiosyncratic states, the mean error remains below 0.08%, the 90th percentile below 0.17%, and the 99.9th percentile below 0.4%.[54] We conclude that the household policies are approximated to a high accuracy.

5.5.3 Stock price

The left panel of figure 6 shows the distribution of discrepancies between the equity price predicted by the neural network and the market clearing equity price obtained using the EGM-Newton-Raphson algorithm. The mean error of the price prediction is 0.17%, the 90th percentile 0.35% and the 99.9th percentile is 0.80%. The right panel of figure 6 shows the price predictions by the neural network together with the computed market clearing prices in a scatter plot against aggregate dividends. Although the accuracy of the price function is slightly lower than the accuracy we achieved for the policy functions of firms and households, it remains high. The resulting errors in asset demand are less than the price errors and less than 0.1% over the entire simulated ergodic set.

[53] Where we evaluate the expectations over $t+1$ objects using policies and prices encoded by neural networks.

[54] For a fixed idiosyncratic asset holding and a fixed idiosyncratic productivity level, the mean and the percentiles are computed across 1024 aggregate states drawn from the ergodic distribution of the economy.

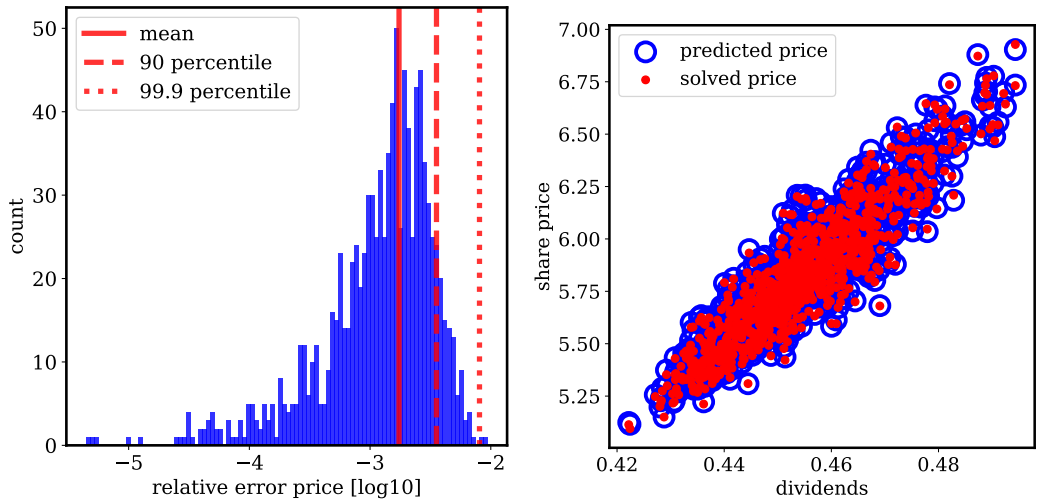


Figure 6: Prediction errors for the price network. The left panel shows the distribution of relative prediction errors across 1024 aggregate states drawn from the ergodic distribution of the economy. The right panel shows the network predictions (blue circles) and the solved market clearing prices (red dots).

5.6 Inspecting the learned equilibrium policies

5.6.1 Firms

We turn to inspecting the equilibrium policies. The left panel of figure 7 shows the density of firms across the capital grid for different simulated aggregate states. We can see that the distribution of firms across capital varies substantially. Due to the decreasing returns to scale production function, together with capital adjustment costs, this variation directly affects aggregate wages and dividends (see, *e.g.*, equation (58)), such that the households' problem depends crucially on the firm distribution.

The middle panel in figure 7 shows the dividend policies of firms for different capital levels (horizontal axis) and different idiosyncratic productivity levels (represented by different colors). The shaded lines show the policies for several randomly selected aggregate states, and the solid lines show the mean taken over aggregate states. Each of the policy functions shows two kinks: at the first kink, the non-negativity constraint on dividends ceases to bind. At the second kink, firms switch from adjusting capital upward to adjusting it downward. The exact location of the kinks depends on the firm's productivity level and the aggregate state of the economy.

Similarly, the right panel shows the policy functions for the KKT multiplier on the dividend constraint. Perhaps surprisingly, the multiplier for very low values of capital is larger for the low productivity values than for the high productivity values. This is the case because, while the firms are constrained for all three productivity levels, the high productivity firms are able to choose a higher capital level for the next period. As guaranteed by the monotonicity preserving neural network architecture, the shaded lines illustrate that the KKT multiplier is weakly decreasing and bounded from below by zero.

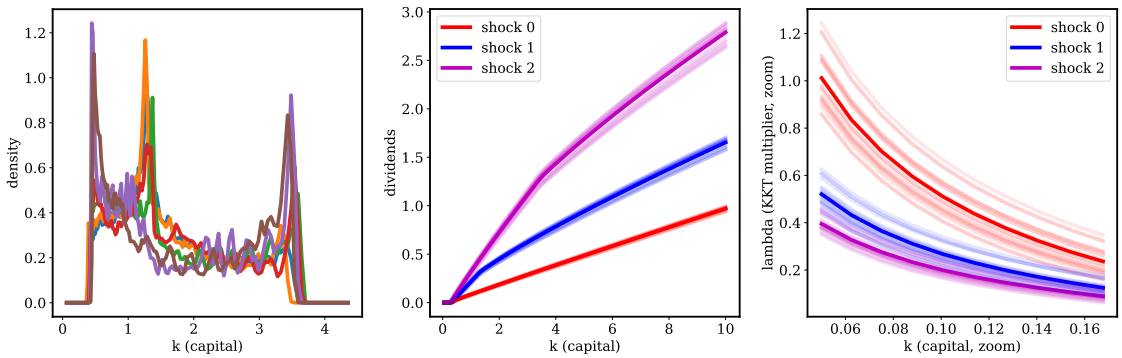


Figure 7: Left panel: density of firms across the capital grid for different simulated aggregate states. Middle panel: dividends policy for firms over the capital grid (horizontal axis) and across different productivity levels (different colors). The shaded lines show the policies for different aggregate states, and the solid lines show the average taken across the sample of aggregate states. Right panel: firms' KKT multiplier associated with the non-negative dividend constraint over the different capital and productivity levels, as well as across different aggregate states.

5.6.2 Households

Figure 8 shows the household policies predicted by the neural network. Each of the shaded lines corresponds to a different aggregate state, and the solid lines correspond to the mean, averaged over aggregate states, for a specific idiosyncratic productivity and wealth level.

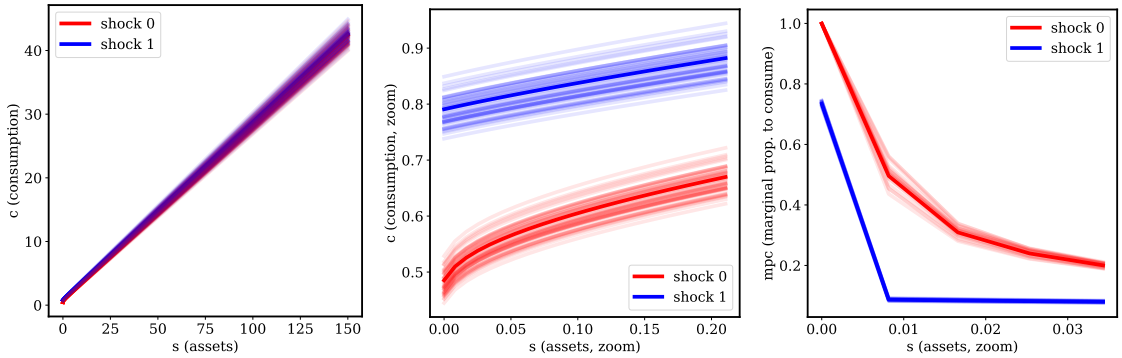


Figure 8: Left and middle panel: consumption function of households for their idiosyncratic productivity level (in different colors) and for different idiosyncratic asset holdings (horizontal axis). The shaded lines show the consumption policies for several different aggregate states, and the solid lines show the mean across aggregate states. Right panel: marginal propensity to consume for different productivity and wealth levels, as well as for several aggregate states. For the asset level $s = 0$, the plot shows the consumption share out of cash at hand.

As figure 8 shows, consumption varies substantially between different aggregate states. We can also see that, for each of the aggregate states, consumption is increasing and concave in the wealth of households. Our shape-preserving neural network architecture ensures that this economically important feature of the consumption functions is guaranteed. We achieve this using the three-

step procedure described above. The households' marginal propensities to consume are constructed from the predictions of the neural network, such that the marginal propensity to consume, for a given aggregate state and idiosyncratic productivity level, is always decreasing in cash-at-hand, as well as bounded between 1 and 0 (see the right panel in figure 8). This ensures on the one hand the monotonicity and concavity of the consumption function, and on the other hand facilitates the precise prediction of the marginal propensity to consume (and hence consumption) of borrowing constraint households.

6 Application to heterogeneous households with discrete and continuous choices

In this section, we illustrate that our method is also applicable to economies where individual agents face a mix of discrete and continuous choices. The combination of discrete and continuous choices poses an additional challenge because it often generates convex-concave value functions.[55] As a result, the Euler equations characterizing the continuous choices generally have multiple solutions in the areas of the state space that feature a locally convex value function. Hence, Euler equations alone are not *sufficient* for characterizing agents' behavior. Additionally, the presence of discrete choices can lead to discontinuous policy functions, which can therefore only be approximated with flexible function approximators. Nevertheless, the sequence-space approach, coupled with shape-preserving operator learning allows us to accurately solve a general equilibrium life-cycle model featuring aggregate shocks, uninsurable idiosyncratic risk, and a discrete early retirement decision.[56]

6.1 Model

The economy is inhabited by finitely lived households, who face an *optional* early retirement decision. During their lifetime, households' labor endowment is subject to idiosyncratic risk. Households can self-insure by investing in claims to physical capital subject to a no short-sale constraint. The only source of aggregate risk is a three-state Markov process that determines the depreciation rate of capital.

6.1.1 Technology

A representative firm rents capital K_t and efficient units of labor L_t on competitive spot markets and pays a rental rate r_t^K for capital and wage w_t per efficiency unit of labor. The firm produces

[55] See, for example, Fella (2014); Iskhakov et al. (2017); Druedahl and Jørgensen (2017), and Druedahl (2021).

[56] See Bardóczy (2020) for an application of Sequence-Space Jacobians, i.e. the seminal perturbation method developed in Auclert et al. (2021), to models with discrete and continuous choices.

output Y_t using a Cobb-Douglas production technology

$$Y_t = K_t^{\alpha} L_t^{1-\alpha} \tag{113}$$

$$r_t^K = \alpha K_t^{\alpha-1} L_t^{1-\alpha} \tag{114}$$

$$w_t = (1-\alpha) K_t^{\alpha} L_t^{-\alpha}. \tag{115}$$

6.1.2 Aggregate risk

The exogenous part of the aggregate state follows a three-state Markov process $z_t \in \mathcal{Z} := \{z_0, z_1, z_2\}$ with transition matrix $\Pi^z \in \mathbb{R}^{3\times 3}$. The state z_t determines the depreciation rate of capital

$$\delta_t = \delta(z_t). \tag{116}$$

The return on capital carried from t to $t{+}1$ equals $R_{t+1} = 1 - \delta_{t+1} + r_{t+1}^K$. We interpret z_2 as a disaster state with substantially higher depreciation rate. States z_0 and z_1 are "normal" regimes with the same current depreciation rate but different transition probabilities, so they differ in *expected* future depreciation even when δ_t is identical.

6.1.3 Demographics

Households live for $H = 10$ periods. Age is indexed by $h \in \mathcal{H} := \{0, 1, \ldots, H-1\}$. Households work for the first six periods,

$$\mathcal{H}^{\text{work}} := \{0, 1, 2, 3, 4, 5\}, \tag{117}$$

make an *optional* retirement decision at age

$$h \in \mathcal{H}^{\text{optional}} := \{6\}, \tag{118}$$

and are all retired for the last three periods of their life

$$h \in \mathcal{H}^{\text{retired}} := \{7, 8, 9\}. \tag{119}$$

We assume a stationary population with equal mass in each cohort. Households die with certainty at the end of the final period of their life. Bequests left by the dying cohort become the initial assets of newborns in the next period.

6.1.4 Idiosyncratic labor productivity

Each household draws an idiosyncratic productivity state $e_t \in \mathcal{E}^{\text{grid}} = \{e_0, \ldots, e_{N_e-1}\}$ that follows a time-homogeneous Markov chain with transition matrix Π^e. Efficiency units of labor supplied by a working household of age h are $n^h e_t$, where n^h follows a deterministic life-cycle profile. Idiosyncratic productivity is dynastic: newborns inherit the family productivity state (on top of the remaining wealth at death), which continues to evolve according to Π^e across generations.

6.1.5 Preferences, optional retirement, and bequests

Households have CRRA utility over consumption,

$$u(c) = \frac{c^{1-\gamma} - 1}{1-\gamma}. \tag{120}$$

At the optional retirement age $h = 6$, households choose whether to work ($\ell = 1$) or retire early ($\ell = 0$). Working in the optional period carries a deterministic utility cost $\epsilon > 0$. To obtain non-degenerate choice probabilities, we add i.i.d. Type-I extreme value taste shocks ε_i to the discrete alternatives $i \in \{\text{work}, \text{ret}\}$. Let $\sigma > 0$ denote the scale parameter. Finally, in the last period of life ($h = 9$) households have a warm-glow bequest motive over the assets passed on to newborns in the next period. Let $\psi^b > 0$ denote the bequest strength, $\nu > 0$ its curvature, and $\bar{b} > 0$ a shifter:

$$b(k_{t+1}) = \psi^b \frac{\left(k_{t+1} + \bar{b}\right)^{1-\nu} - 1}{1-\nu}. \tag{121}$$

6.1.6 Asset markets and non-market income

Capital is the only asset. Households face a borrowing constraint

$$k_{t+1} \geq \underline{k}. \tag{122}$$

Retired households (including early retirees) receive a constant income y^{home} from home production.

6.1.7 Household problem

Let the individual labor choice be $\ell_t \in \{0, 1\}$ (to avoid confusion with aggregate labor L_t). Outside the optional retirement age, ℓ_t is exogenously fixed:

$$\ell_t = 1 \text{ for } h \in \mathcal{H}^{\text{work}}, \qquad \ell_t \in \{0,1\} \text{ for } h = 6, \qquad \ell_t = 0 \text{ for } h \in \mathcal{H}^{\text{retired}}. \tag{123}$$

The household budget constraint is

$$c_t = R_t k_t + y_t(h, e_t, \ell_t) - k_{t+1}, \tag{124}$$

where income $y_t(h, e_t, l)$ is given by

$$y_t(h, e, \ell) := \begin{cases} w_t \, n^h e, & \text{if } \ell = 1, \\ y^{\text{home}}, & \text{if } \ell = 0. \end{cases} \tag{125}$$

It is helpful to define the *savings-choice-specific* continuation value

$$W_t(h, e_t, k_{t+1}) = \begin{cases} \mathbb{E}_t[V_{t+1}(h+1, e_{t+1}, k_{t+1})], & \text{for } h < 9, \\ b(k_{t+1}), & \text{for } h = 9, \end{cases} \tag{126}$$

where $\mathbb{E}_t[\cdot]$ is taken with respect to the idiosyncratic transition Π^e and the aggregate regime transition Π^z, conditional on current states and where $V(\cdot)$ denotes the value function. The savings-choice-specific continuation values, $W_t(h, e_t, k_{t+1})$, are the equilibrium objects, which we approximate with neural networks.

Continuous savings problem (all ages except $h = 6$) For $h \in \{\mathcal{H}^{\text{work}} \cup \mathcal{H}^{\text{retired}}\}$, the Bellman equation is

$$V_t(h, e_t, k_t) = \max_{k_{t+1} \geq \underline{k}} \{u(c_t) + \beta W_t(h, e_t, k_{t+1})\}, \tag{127}$$

subject to the borrowing constraint (122) and the budget constraint (124).

Optional retirement at $h = 6$ At age $h = 6$ we work with discrete-choice-specific value functions excluding taste shocks. Let $\tilde{V}_t^{\text{work}}(e_t, k_t)$ and $\tilde{V}_t^{\text{ret}}(e_t, k_t)$ solve

$$\tilde{V}_t^{\text{work}}(e_t, k_t) = \max_{k_{t+1} \geq \underline{k}} \left\{u(c_t^{\text{work}}) - \epsilon + \beta W_t(6, e_t, k_{t+1})\right\}, \tag{128}$$

$$\tilde{V}_t^{\text{ret}}(e_t, k_t) = \max_{k_{t+1} \geq \underline{k}} \left\{u(c_t^{\text{ret}}) + \beta W_t(6, e_t, k_{t+1})\right\}, \tag{129}$$

where c_t^{work} uses $\ell = 1$ (labor income $w_t n^6 e_t$) and c_t^{ret} uses $\ell = 0$ (home income y^{home}) in the budget constraint (124). Since both choices lead to retirement from age $h = 7$ onward, the conditional continuation term $W_t(6, e_t, k_{t+1})$ is common across the two alternatives. With i.i.d. Type-I extreme value taste shocks of scale σ, the ex-ante value is

$$V_t(6, e_t, k_t) = \sigma \log\Big(\exp(\tilde{V}_t^{\text{work}}(e_t, k_t)/\sigma) + \exp(\tilde{V}_t^{\text{ret}}(e_t, k_t)/\sigma)\Big), \tag{130}$$

and the conditional probability of working is

$$p_t^{\text{work}}(e_t, k_t) = \frac{\exp\big(\tilde{V}_t^{\text{work}}(e_t, k_t)/\sigma\big)}{\exp\big(\tilde{V}_t^{\text{work}}(e_t, k_t)/\sigma\big) + \exp\big(\tilde{V}_t^{\text{ret}}(e_t, k_t)/\sigma\big)}. \tag{131}$$

In the limit $\sigma \to 0$, this expression converges to the deterministic discrete choice.

6.1.8 Equilibrium

State space The aggregate state is

$$\mathbf{X}_t^{\text{agg. state}} = [z_t, \mu_t(h, e, k)], \tag{132}$$

where z_t is the aggregate depreciation regime and $\mu_t(h, e, k)$ is the cross-sectional distribution of households over age, idiosyncratic productivity and assets.

Aggregation and prices The aggregate capital supply is given by

$$K_t = \sum_{h \in \mathcal{H}} \sum_{e \in \mathcal{E}^{\text{grid}}} \int_k k_t(h,e,k)\ \mu_t(h,e,k) \mathrm{d}k. \tag{133}$$

Efficient units of labor supplied by an individual of age h are $\ell_t(h,e,k)\, n^h e$, where

$$\ell_t(h,e,k) = \begin{cases} 1, & \text{for } h \in \mathcal{H}^{\text{work}}, \\ p_t^{\text{work}}(e,k), & \text{for } h = 6, \\ 0, & \text{for } h \in \mathcal{H}^{\text{retired}}. \end{cases} \tag{134}$$

The aggregate labor supply is given by

$$L_t = \sum_{h \in \mathcal{H}} \sum_{e \in \mathcal{E}^{\text{grid}}} \int_k \ell_t(h,e,k) n^h\, e\ \mu_t(h,e,k) \mathrm{d}k. \tag{135}$$

It is useful to summarize the optional cohort's labor supply by the fraction of its *maximum* possible efficiency units supplied:

$$L_t^{6,\max} := \sum_{e \in \mathcal{E}^{\text{grid}}} \int_k n^6\, e\ \mu_t(6,e,k) \mathrm{d}k, \tag{136}$$

$$L_t^{6} := \sum_{e \in \mathcal{E}^{\text{grid}}} \int_k p_t^{\text{work}}(e,k)\, n^6\, e\ \mu_t(6,e,k) \mathrm{d}k, \tag{137}$$

$$\lambda_t := \frac{L_t^6}{L_t^{6,\max}} \in [0,1]. \tag{138}$$

By construction, $\lambda_t = 1$ if no one retires early at $h = 6$, and $\lambda_t = 0$ if everyone retires early. Using λ_t, aggregate labor can be written as

$$L_t = \sum_{h \in \mathcal{H}^{\text{work}}} \sum_{e \in \mathcal{E}^{\text{grid}}} \int_k n^h\, e\ \mu_t(h,e,k) \mathrm{d}k.\ +\ \lambda_t\, L_t^{6,\max}. \tag{139}$$

Given (K_t, L_t), prices can be computed in closed form using the firm's first-order conditions (114) and (115). Market clearing conditions consist of: (i) capital-market clearing, where the representative firm rents the aggregate household capital stock K_t; (ii) labor-market clearing, where firm labor demand equals aggregate labor supply L_t; and (iii) goods-market clearing. Given that (i) and (ii) hold, (iii) is implied by Walras' law.

Functional rational expectations equilibrium A functional rational expectations equilibrium consists of (i) cohort- and productivity-specific savings and consumption functions for all ages, (ii) choice-specific savings and value functions at the optional retirement age $h = 6$ that imply a logit working probability p_t^{work}, and (iii) an induced law of motion for μ_t such that households' choices are optimal given prices implied by (K_t, L_t) and the evolution of aggregate regimes, and markets clear.

6.2 Parameterization

One model period corresponds to six calendar years. We set $\gamma = 2$ and $\beta = 0.95^6 \approx 0.735$. Idiosyncratic productivity takes three values $e \in \mathcal{E}^{\text{grid}} = \{0.5, 1.0, 1.5\}$ with transition matrix

$$\Pi^e = \begin{bmatrix} 0.70 & 0.20 & 0.10 \\ 0.15 & 0.70 & 0.15 \\ 0.10 & 0.20 & 0.70 \end{bmatrix}. \tag{140}$$

We impose $\underline{k} = 0$ and approximate the asset space with an evenly spaced grid of $N_k = 300$ points on $[0, 5.0]$. We set $\alpha = \frac{1}{3}$. Depreciation is regime-dependent with

$$\delta(z_0) = \delta(z_1) = \delta_n, \qquad \delta(z_2) = \delta_d, \tag{141}$$

where $\delta_n = 1 - (1 - 0.05)^6 \approx 0.265$ and $\delta_d = 1.5\,\delta_n \approx 0.397$. The Markov transition matrix for z_t is

$$\Pi^z = \begin{bmatrix} 0.90 & 0.07 & 0.03 \\ 0.70 & 0.20 & 0.10 \\ 0.50 & 0.25 & 0.25 \end{bmatrix}. \tag{142}$$

The life-cycle component n^h is set to

$$(n^0, \ldots, n^6) = (0.53,\ 0.83,\ 1.04,\ 1.17,\ 1.21,\ 1.17,\ 1.04), \tag{143}$$

and $n^h = 0$ for $h \in \{7, 8, 9\}$. The deterministic utility cost of working in the optional period is set to $\epsilon = 1.68$ and the logit smoothing parameter to $\sigma = 0.08$. Retirees receive $y^{\text{home}} = 0.1$ each period. For bequests, we set $\psi^b = 0.1$, $\nu = 2$, and $\bar{b} = 0.1$. The parameters are summarized in table A7.

6.3 Implementation

We rely on the monotone operator framework introduced in section 5.3.1. In contrast to the previous applications, we chose to approximate the choice-specific continuation values. Given prices and the predicted continuation values, savings policies and the model-implied work probabilities are computed via generalized EGM (with an upper-envelope step) and the logit choice rule.

6.3.1 State representation

The endogenous state is the cross-sectional distribution μ_t over age $h \in \mathcal{H}$, idiosyncratic productivity $e \in \mathcal{E}^{\text{grid}}$, and assets $k \in \mathcal{K}^{\text{grid}}$. We represent μ_t as a histogram on the tensor grid

$$\mathcal{G}^{\text{grid}} := \mathcal{H} \times \mathcal{E}^{\text{grid}} \times \mathcal{K}^{\text{grid}},$$

with $H = 10$, $N_e = 3$, and $N_k = 200$ grid points. We write $\mu_t(h, e_i, k_j)$ for the mass in cell (h, e_i, k_j). Following the sequence-space approach, the neural operator takes a truncated history of regime

realizations as input

$$\mathbf{X}_t^{\text{seq},T} = (z_t, z_{t-1}, \ldots, z_{t-T+1}).$$

6.3.2 Household policies with operator learning

Let $\mathcal{N}_{\boldsymbol{\rho}}$ denote the neural operator with parameters $\boldsymbol{\rho}$. Given a history $\mathbf{X}_t^{\text{seq},T}$, the operator outputs two objects: (i) grids of continuation values for all cohorts except the terminal cohort, and (ii) a scalar guess for optional-cohort labor supply.

Continuation values We predict the savings-choice-specific continuation value $W_t(h, e, k')$ on the fixed asset grid for $h \in \{0, \ldots, 8\}$:

$$\mathcal{N}_{\boldsymbol{\rho}}^{W}(\mathbf{X}_t^{\text{seq},T}) = \hat{\mathcal{W}}_t = \{\hat{W}_t(h, e_i, k_j) \mid h \leq 8,\ e_i \in \mathcal{E}^{\text{grid}},\ k_j \in \mathcal{K}^{\text{grid}}\}. \tag{144}$$

For the terminal cohort, the continuation value is given by the bequest function, $W_t(9, e, k') = b(k')$.

To enforce monotonicity in assets, we parameterize each $\hat{W}_t(h, e, \cdot)$ by a boundary value and non-negative increments. Specifically, the network outputs $\hat{W}_t(h, e, k_0)$ and raw increments $\tilde{\Delta}_t(h, e, k_j)$ for $j = 1, \ldots, N_k - 1$, sets $\Delta_t(h, e, k_j) = \text{softplus}(\tilde{\Delta}_t(h, e, k_j))$, and constructs

$$\hat{W}_t(h, e, k_j) = \hat{W}_t(h, e, k_0) + \sum_{m=1}^{j} \Delta_t(h, e, k_m),$$

which guarantees $\hat{W}_t(h, e, k)$ is increasing in k.

Optional-cohort labor-supply We also output a scalar

$$\mathcal{N}_{\boldsymbol{\rho}}^{\lambda}(\mathbf{X}_t^{\text{seq},T}) = \hat{\lambda}_t \in (0, 1), \tag{145}$$

interpreted as the fraction of the optional cohort's *maximum* possible efficiency units supplied (so $\lambda_t = 1$ corresponds to no early retirement and $\lambda_t = 0$ to universal early retirement). The scalar prediction $\hat{\lambda}_t$[57] is used only to initialize Newton's method for labor-market clearing (Step 0 in Section 6.3.4); the realized λ_t is computed from the household discrete-choice problem given equilibrium prices.

From continuation values to policies Given (K_t, z_t), a candidate labor input L determines prices (w_t, r_t^K). Using these prices and the predicted continuation values $\hat{\mathcal{W}}_t$, we compute savings and consumption policies for all cohorts via generalized EGM. For worker cohorts we apply an upper-envelope step to handle potential nonconcavities. At $h = 6$, we solve the two conditional saving problems (work/retire), compute the corresponding choice-specific value functions, and obtain the model-implied work probabilities $p_t^{\text{work}}(e, k)$ from the logit formula.

[57] In our implementation we enforce $\hat{\lambda}_t \in (0, 1)$ by applying a sigmoid activation in the output layer of $\mathcal{N}_{\boldsymbol{\rho}}^{\lambda}$.

6.3.3 Loss function

Let the network outputs be the continuation-value grids $\hat{\mathcal{W}}_t$ and the scalar $\hat{\lambda}_t$. Given supervised learning targets $(\mathcal{W}_t^{\text{tar}}, \lambda_t^{\text{tar}})$ constructed by the equilibrium routine described in Section 6.4, we minimize mean squared *relative* errors on economically interpretable components: (i) the fraction of optional-cohort labor-supply, (ii) boundary values of the continuation-value, and (iii) slope of the continuation-value on the asset grid.

For a small $\varepsilon_{\text{loss}} > 0$,[58] define

$$\text{err}_\lambda = \frac{\lambda_t^{\text{tar}} - \hat{\lambda}_t}{\varepsilon_{\text{loss}} + \lambda_t^{\text{tar}}}, \tag{146}$$

$$\text{err}_{W,0}(h,e) = \frac{W_t^{\text{tar}}(h,e,k_1) - \hat{W}_t(h,e,k_1)}{\varepsilon_{\text{loss}} + \left|W_t^{\text{tar}}(h,e,k_1)\right|}, \tag{147}$$

$$\text{err}_{W,\Delta}(h,e,j) = \frac{\Delta W_t^{\text{tar}}(h,e,k_j) - \Delta \hat{W}_t(h,e,k_j)}{\varepsilon_{\text{loss}} + \left|\Delta W_t^{\text{tar}}(h,e,k_j)\right|}, \tag{148}$$

where $\Delta W(h,e,k_j) = W(h,e,k_j) - W(h,e,k_{j-1})$ and $\Delta \hat{W}_t(h,e,k_j) := \hat{W}_t(h,e,k_j) - \hat{W}_t(h,e,k_{j-1})$. The loss is the mean of err_λ^2, $\text{err}_{W,0}^2$, and $\text{err}_{W,\Delta}^2$ over grid indices and sampled aggregate state-history pairs:

$$\mathcal{L} = \frac{1}{B}\sum_{b=1}^{B} \text{err}_{\lambda,b}^2 + \frac{1}{B}\sum_{b=1}^{B} \frac{1}{|\mathcal{I}_{W,0}|} \sum_{(h,e)\in\mathcal{I}_{W,0}} \text{err}_{W,0,b}(h,e)^2 + \frac{1}{B}\sum_{b=1}^{B} \frac{1}{|\mathcal{I}_{W,\Delta}|} \sum_{(h,e,j)\in\mathcal{I}_{W,\Delta}} \text{err}_{W,\Delta,b}(h,e,j)^2, \tag{149}$$

where $\mathcal{I}_{W,0} = \{(h,e) : h \le 8, e \in \mathcal{E}^{\text{grid}}\}$ and $\mathcal{I}_{W,\Delta} = \{(h,e,j) : h \le 8, e \in \mathcal{E}^{\text{grid}}, j = 2,\dots,N_k\}$. In the baseline implementation we set $\varepsilon_{\text{loss}} = 10^{-2}$.

6.3.4 Forward simulation

This subsection describes the forward simulation step that maps the current aggregate state (z_t, μ_t) and regime history $\mathbf{X}_t^{\text{seq},T}$ into a successor state-history pair used for training, and (as part of the same routine) constructs the objects needed to form continuation-value targets.

Step 0: equilibrium policies and labor market clearing For each state-history pair $(z_t, \mu_t, \mathbf{X}_t^{\text{seq},T})$ we first solve for the within-period general equilibrium objects and policy functions. Given the current histogram μ_t, aggregate capital is predetermined. Aggregate labor L_t is unknown because (i) wages and the rental rate follow from (K_t, L_t, z_t) and (ii) the optional-cohort participation probability depends on wages and thus on L_t. We therefore solve for a scalar fixed point L satisfying:

$$ED(L) = L - \mathcal{L}^{hh}(L) = 0,$$

[58] Since the slope of the continuation value can become close to zero during training, we use $\varepsilon_{\text{loss}} > 0$ in order to define approximately relative errors.

where $\mathcal{L}^{hh}(L)$ is the aggregate labor implied by household policies when prices are computed for aggregates (K_t, L, z_t). In our implementation we take a small, fixed number of Newton steps and obtain $ED'(L)$ by automatic differentiation. A convenient initial guess for L is constructed as $L^{\text{guess}} = L^{\text{work}} + \hat{\lambda}_t\, L^{6,\max}$, where L^{work} are efficiency units supplied by the non-optional working cohorts, $L^{6,\max}$ is the optional cohort's maximum possible efficiency units (i.e. if nobody retires early), and $\hat{\lambda}_t \in (0,1)$ is the network-predicted fraction of that maximum.

Each evaluation of $ED(L)$ calls the generalized EGM to solve the household problems at all cohorts given (i) prices implied by (K_t, L, z_t) and (ii) the network-predicted continuation objects for the relevant cohorts. This yields: (a) savings policies for worker cohorts $\{g_t^h(e,k)\}_{h\in\mathcal{H}^{\text{work}}}$, (b) the two conditional savings policies at the optional cohort, $g_t^{6,\text{work}}(e,k)$ and $g_t^{6,\text{ret}}(e,k)$, together with the logit work probabilities $p_t^{\text{work}}(e,k)$, and (c) savings policies for retired cohorts including the terminal bequest policy. The converged Newton iterate provides the equilibrium L_t and the policy bundle that is used in the distribution transition below.

Step 1: simulating the distribution forward Let $\mathcal{K}^{\text{grid}} = \{k_1 < \cdots < k_{N_k}\}$ and $\mathcal{E}^{\text{grid}} = \{e_1, \ldots, e_{N_e}\}$. Write $\mu_t(h,i,j)$ for the histogram mass at age h, productivity e_i, and assets k_j. For a cohort-specific policy $g(e_i, k_j)$, the implied next assets $k' = g(e_i, k_j)$ typically lie off-grid. Following the non-stochastic simulation method introduced by Young (2010), we project mass to the fixed grid by splitting it linearly between the two bracketing grid points. Let $j^-(i,j)$ and $j^+(i,j)$ denote indices such that $k_{j^-} \le k' \le k_{j^+}$ and define weights

$$\omega^-(k') = \frac{k_{j^+} - k'}{k_{j^+} - k_{j^-}}, \qquad \omega^+(k') = 1 - \omega^-(k').$$

Idiosyncratic productivity transitions are applied exactly using the transition matrix Π^e. Define the corresponding transport operator $\mathcal{T}$ acting on a mass array $m(i,j)$:

$$\big[\mathcal{T}(g,m)\big](i',j') = \sum_{i=1}^{N_e}\sum_{j=1}^{N_k} m(i,j)\,\Pi^e_{ii'}\Big(\omega^-(g(i,j))\,\mathbf{1}\{j' = j^-(i,j)\} + \omega^+(g(i,j))\,\mathbf{1}\{j' = j^+(i,j)\}\Big). \tag{150}$$

This step is deterministic and conserves mass by construction.

Step 2: cohort aging, optional retirement split, and newborns Using (150), the histogram update is given by

$$\mu_{t+1}(0,\cdot,\cdot) = \mathcal{T}\big(g_t^9,\, \mu_t(9,\cdot,\cdot)\big)\,, \tag{151}$$

$$\mu_{t+1}(h+1,\cdot,\cdot) = \mathcal{T}\big(g_t^h,\, \mu_t(h,\cdot,\cdot)\big)\,, \qquad h \in \{0,1,2,3,4,5\}, \tag{152}$$

$$\mu_{t+1}(7,\cdot,\cdot) = \mathcal{T}\Big(g_t^{6,\text{work}},\, p_t^{\text{work}} \odot \mu_t(6,\cdot,\cdot)\Big) + \mathcal{T}\Big(g_t^{6,\text{ret}},\, (1 - p_t^{\text{work}}) \odot \mu_t(6,\cdot,\cdot)\Big)\,, \tag{153}$$

$$\mu_{t+1}(h+1,\cdot,\cdot) = \mathcal{T}\big(g_t^h,\, \mu_t(h,\cdot,\cdot)\big)\,, \qquad h \in \{7,8\}. \tag{154}$$

Equation (153) is the only place where the discrete choice enters the forward step: the optional cohort's mass is split according to the EGM-implied logit probability $p_t^{\text{work}}(e,k)$, and each part follows its corresponding conditional savings policy. Equation (151) maps the assets of the terminal cohort to the newborn distribution via the bequest policy; since (150) applies Π^e, productivity transitions are applied uniformly in all cohort transitions, including between generations.

Step 3: selection of the next state-history pair The distribution update (151)–(154) does not depend on the realized z_{t+1}; only the regime component changes. Hence we form candidate successor states for all regimes $z' \in \{z_0, z_1, z_2\}$ as

$$\mathbf{X}_{t+1}^{\text{agg. state}}(z') = (z', \mu_{t+1}), \qquad \mathbf{X}_{t+1}^{\text{seq},T}(z') = (z', z_t, z_{t-1}, \ldots, z_{t-T+2}).$$

To evolve the simulated training state-history pairs, we draw a realized regime $z_{t+1} \sim \Pi^z(z_t, \cdot)$ and select the corresponding $(\mathbf{X}_{t+1}^{\text{agg. state}}(z_{t+1}), \mathbf{X}_{t+1}^{\text{seq},T}(z_{t+1}))$.

6.4 Training

6.4.1 Outline of the training procedure

We train the neural network with supervised learning on simulated data. Each training point consists of (i) an aggregate regime z_t, (ii) the cross-sectional distribution μ_t on the tensor grid (h, e, k), and (iii) a truncated regime history $\mathbf{X}_t^{\text{seq},T} = (z_t, z_{t-1}, \ldots, z_{t-T+1})$ that serves as the network input. Given $\mathbf{X}_t^{\text{seq},T}$, the network outputs (a) continuation-value arrays for all cohorts except the terminal cohort, parameterized in a monotone way over the asset grid (boundary plus positive increments), and (b) a scalar prediction $\hat{\lambda}_t \in (0,1)$ for the optional cohort's effective participation rate. The scalar $\hat{\lambda}_t$ is used only to initialize the labor-market-clearing routine. The training procedure is organized in three stages.

Stage 1: deterministic steady state We first compute a deterministic steady state with constant depreciation and no regime switching. We iterate on a guess for (K, L), compute prices, solve the life-cycle problem via EGM (including the optional retirement decision), and simulate the stationary distribution. The resulting steady-state policy, value objects and stationary distribution are stored and used to initialize (i) the simulated training data and (ii) the pretraining stage below.

Stage 2: stochastic pretraining with steady-state continuation We next introduce aggregate risk, but when solving the household problems we hold the *next-period* continuation objects fixed at their steady-state values. Operationally, the EGM step uses steady-state continuation values on the right-hand side of the Euler equation and in the discrete-choice log-sum, while prices and the current distribution are allowed to vary with the simulated history. This pretraining stage stabilizes optimization and produces a diversified cloud of $(z_t, \mu_t, \mathbf{X}_t^{\text{seq},T})$ before turning on full recursion.

Stage 3: full general equilibrium Finally, we train in the full model in which continuation objects depend on the operator itself. To stabilize the training, we use a target network, whose pa-

rameters are updated more slowly. Let $\boldsymbol{\rho}$ denote the online parameters and $\bar{\boldsymbol{\rho}}$ the target parameters, updated by

$$\bar{\boldsymbol{\rho}} \leftarrow \tau\, \boldsymbol{\rho} + (1-\tau)\, \bar{\boldsymbol{\rho}}.$$

At each iteration we take one supervised-learning step on a batch of size B. For each state-history pair $(z_t, \mu_t, \mathbf{X}_t^{\text{seq},T})$, one iteration consists of:

1. *Online prediction at t.* Evaluate the online network to obtain predictions $\mathcal{N}_{\boldsymbol{\rho}}(\mathbf{X}_t^{\text{seq},T})$, including the continuation-value arrays and the scalar participation guess $\hat{\lambda}_t$.

2. *Target prediction at candidate $t{+}1$ histories.* Construct the candidate next histories $\{\mathbf{X}_{t+1}^{\text{seq},T}(z')\}_{z'\in\mathcal{Z}}$ by prepending each possible next regime z' and shifting the history vector. Evaluate the *target* network on these candidate histories to obtain $\mathcal{N}_{\bar{\boldsymbol{\rho}}}(\mathbf{X}_{t+1}^{\text{seq},T}(z'))$ for all $z' \in \mathcal{Z}$.

3. *Current equilibrium with market clearing.* Given (z_t, μ_t) and the online predictions $\mathcal{N}_{\boldsymbol{\rho}}(\mathbf{X}_t^{\text{seq},T})$, we clear the labor market by solving

$$ED(L_t) = L_t - \mathcal{L}^{hh}(L_t) = 0$$

   with a small, fixed number of Newton–Raphson iterations. Here $\mathcal{L}^{hh}(L)$ denotes aggregate labor implied by household policies under prices computed at (K_t, L, z_t). Each evaluation of $ED(L)$ solves the household problems for all cohorts via EGM using the online network's continuation objects. The derivative $ED'(L)$ used in Newton's method is obtained by automatic differentiation.

   This step delivers equilibrium prices and the full policy bundle at t, including (i) the worker savings policies, (ii) the optional-cohort work probabilities and the two conditional savings rules (work or retire), and (iii) the retirement and terminal bequest policies, as well as the *realized* optional-cohort participation rate λ_t.

4. *Forward simulation of the distribution.* Using the current equilibrium policy bundle from step (3), we update the histogram deterministically to obtain μ_{t+1}, following Young (2010). This produces the aggregate states in the next period (z', μ_{t+1}) for all $z' \in \mathcal{Z}$.

5. *Next-period equilibria for target construction.* For each possible next regime z', we solve the $t{+}1$ equilibrium as in step (3), but using the *target-network* predictions $\mathcal{N}_{\bar{\boldsymbol{\rho}}}(\mathbf{X}_{t+1}^{\text{seq},T}(z'))$ to evaluate next-period continuation objects in the EGM step. We then form continuation-value targets at t by taking expectations over z_{t+1} and e_{t+1} using the Markov transition matrices.

6. *Loss and parameter updates* We define supervised targets as (i) the continuation-value arrays constructed in step (5) and (ii) the realized optional-cohort effective participation rate λ_t from step (3). The loss is the mean squared relative error between online predictions and the constructed training targets; for continuation values we compute errors on the boundary and on one-step increments (consistent with the monotone parameterization), and for participation we compute a scalar relative error. Gradients are taken only through the loss function, taking

the target values as fixed. We update the online parameters $\boldsymbol{\rho}$ using the Adam optimizer and update the target parameters $\bar{\boldsymbol{\rho}}$ by the EMA rule above.

7. *Evolve the training state-history pairs.* To generate the next on-policy batch, we draw $z_{t+1} \sim \Pi^z(z_t, \cdot)$ and select the corresponding $(z_{t+1}, \mu_{t+1}, \mathbf{X}_{t+1}^{\text{seq},T}(z_{t+1}))$ as the next state history pair.

6.4.2 Hyperparameters

$\mathcal{N}_{\boldsymbol{\rho}}$ is a feed-forward network with GELU activations (three hidden layers of width 1024) and input history length $T = 100$. We train with Adam using a learning rate of 10^{-6} and an EMA parameter $\tau = 10^{-4}$. Each iteration uses a $B = 1024$ state-history pairs and clears the labor market with 5 Newton steps. Relative-error losses use a stabilizer parameter $\varepsilon_{\text{loss}}$ in the denominator to avoid numerical issues when targets are near zero. At the beginning of training we set $\varepsilon_{\text{loss}} = 1.0$. During the training, we reduce $\varepsilon_{\text{loss}}$ first to 0.1, and finally to 10^{-2}.

6.5 Accuracy

To assess the accuracy of our solution, we calculate the distribution of prediction errors of our continuation value network. The continuation values are shown in figure A2 and error statistics are shown in figure A3. Since the continuation values cross zero, relative errors are not meaningful. Because of that, we report absolute errors across idiosyncratic and aggregate states.

6.6 Inspecting the equilibrium

In figure 9 we plot consumption functions for working-age cohorts. Figure 10 shows conditional consumption functions and work probabilities of the cohort facing optional retirement, and figure A4 plots the consumption function of full retirees. The bottom right panel in figure 9 shows the consumption functions for working-age households in the period before the optional early-retirement period. The discrete retirement choice leads to discontinuities in the consumption function of pre-retirement cohorts. The discontinuities propagate backward and slowly dissipate as uncertainty concavifies the continuation values of younger cohorts. The two left panels in figure 10 show that the consumptions conditional on the retirement decision of the optional cohort are smooth. As shown in the right two panels, the discrete retirement choice, however, renders the expected consumption strongly non-linear and non-monotone. This is the source of the non-concavities in the continuation values and strongly non-linear policies in the prior periods. Figure A4 shows the consumption function of retired households. Since retired households do not earn wage, their consumption does not depend on their idiosyncratic productivity state.

7 Conclusion

We introduce a new algorithm for computing global solutions of dynamic stochastic general equilibrium models. Exploiting the ergodic property of a large class of dynamic economies, we rely on the truncated history of aggregate shocks as an approximately sufficient statistic for the aggregate

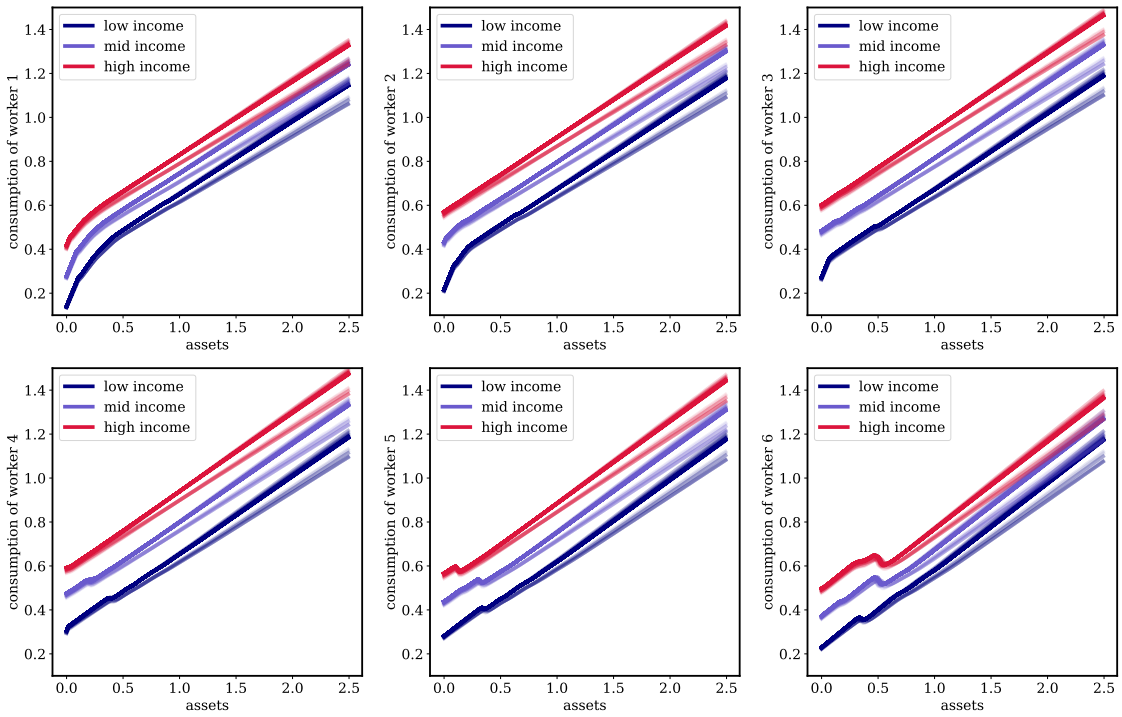


Figure 9: Consumption functions of working-age cohorts. The shaded lines show the policies for different aggregate states, and the solid lines show the average taken across the sample of aggregate states.

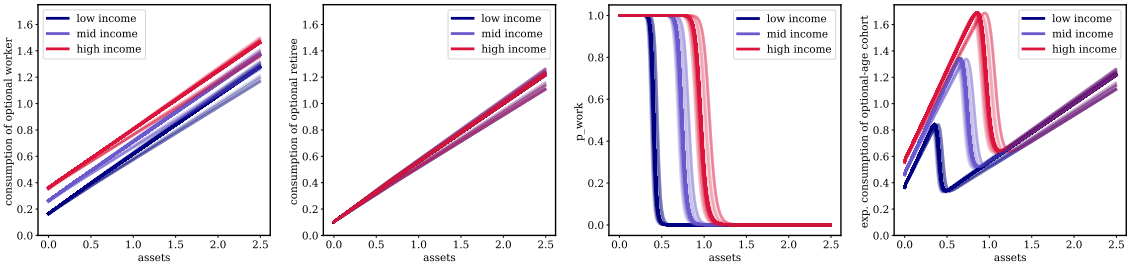


Figure 10: First and second panel: consumption function of the optional retirement cohort conditional on the retirement decision. Third and fourth panel: probability to work in the optional period together with the expected consumption conditional on assets, idiosyncratic productivity and the aggregate history, where the expectation is taken over the taste shock. The shaded lines show different aggregate histories and the solid shows the mean across aggregate histories.

state. Based on this approximation, we use deep neural networks to parameterize the mapping from truncated aggregate shock histories to equilibrium objects of interest. Finally, we train our neural networks to satisfy all the equilibrium conditions along simulated paths of the economy.

We apply our method to three challenging economies, each constructed to feature a distinct difficulty that commonly arises in dynamic equilibrium models. First, we solve an overlapping generations model with 72 age groups, portfolio choice, and multiple sources of aggregate risk. The overlapping generations model showcases that our sequence space approach is applicable even when there is a clear dependence of current equilibrium objects on long histories of shocks.

Second, we consider an economy where a continuum of heterogeneous households trade in a long-lived financial asset, which constitutes a claim to a share in the dividends paid by a continuum of heterogeneous firms, who operate a decreasing returns to scale production technology and face asymmetric adjustment costs on capital. Additionally, the economy is subject to stochastic fluctuations in aggregate productivity and in the level of aggregate and idiosyncratic uncertainty. To solve models that combine aggregate and idiosyncratic risk, we introduce shape-preserving operator learning: we train deep neural networks to predict idiosyncratic policy *functions* as a function of the truncated history of aggregate shocks, such that we can guarantee the predicted policy functions to be monotone or concave in key idiosyncratic state variables, such as household wealth. These guarantees allow us to use the method of endogenous gridpoints and a simple Newton-Raphson algorithm to obtain high-quality training targets for supervised learning of household policies and the market clearing stock price.

As a final example, we solve for an equilibrium in a life-cycle economy with aggregate and idiosyncratic risk where households have an early retirement option. Solving this model presents three key challenges. First, the early retirement option introduces non-convexities into the decision problems of working-age households. Because of that, Euler equations are not sufficient to pin-down consumption and savings behavior of working-age households. Second, the non-convexities make the policy functions discontinuous: expectations of future retirement introduce jumps in the savings function of workers, moreover the location of those jumps depends on the aggregate state. Third, this economy features an especially high-dimensional state space: there is a non-degenerate cross-sectional wealth distribution within each age cohort. We show that our method remains applicable and computationally tractable even in all three environments.

References

Adenbaum, J., Babalievsky, F., and Jungerman, W. (2024). Deep reinforcement learning for economics. Technical report, Working Paper.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684.

Auclert, A., Bardóczy, B., Rognlie, M., and Straub, L. (2021). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408.

Azinovic, M., Cole, H. L., and Kubler, F. (2023). Asset pricing in a low rate environment. Working Paper 31832, National Bureau of Economic Research.

Azinovic, M., Gaegauf, L., and Scheidegger, S. (2022). Deep equilibrium nets. *International Economic Review*, 63(4):1471–1525.

Azinovic, M. and Zemlicka, J. (2024). Intergenerational consequences of rare disasters.

Bardóczy, B. (2020). Spousal insurance and the amplification of business cycles. *Unpublished Manuscript, Northwestern University*.

Bayer, C. and Luetticke, R. (2020). Solving discrete time heterogeneous agent models with aggregate risk and many idiosyncratic states by perturbation. *Quantitative Economics*, 11(4):1253–1288.

Bellman, R. (1961). *Adaptive Control Processes: A Guided Tour*. 'Rand Corporation. Research studies. Princeton University Press.

Bewley, T. (1977). The permanent income hypothesis: A theoretical formulation. *Journal of Economic Theory*, 16(2):252–292.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., and Terry, S. J. (2018). Really uncertain business cycles. *Econometrica*, 86(3):1031–1065.

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

Bretscher, L., Fernández-Villaverde, J., and Scheidegger, S. (2022). Ricardian business cycles. *Available at SSRN 4278274*.

Brock, W. A. and Mirman, L. J. (1972). Optimal economic growth and uncertainty: the discounted case. *Journal of Economic Theory*, 4(3):479–513.

Brumm, J. and Hußmann, J. (2025). Tensor-train approximation for high-dimensional economic models. *Available at SSRN 5989855*.

Brumm, J. and Scheidegger, S. (2017). Using adaptive sparse grids to solve high-dimensional dynamic models. *Econometrica*, 85(5):1575–1612.

Cai, Y. and Judd, K. L. (2012). Dynamic programming with shape-preserving rational spline hermite interpolation. *Economics Letters*, 117(1):161–164.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics letters*, 91(3):312–320.

Carvalho, V. M., Covarrubias, M., and Nuno, G. (2025). Planning against disasters in dynamic production networks. *Working Paper*.

Chien, Y., Cole, H., and Lustig, H. (2011). A multiplier approach to understanding the macro implications of household finance. *The Review of Economic Studies*, 78(1):199–234.

Dobrescu, L. I. and Shanker, A. (2023). *A fast upper envelope scan method for discrete-continuous dynamic programming*. CEPAR, ARC Centre of Excellence in Population Ageing Research.

Druedahl, J. (2021). A guide on solving non-convex consumption-saving models. *Computational Economics*, 58(3):747–775.

Druedahl, J., Huleux, R., and Ropke, J. (2026a). Deep learning solutions of large non-convex life-cycle models. Technical report, University of Copenhagen.

Druedahl, J., Huleux, R., and Ropke, J. (2026b). Fiscal multipliers in a globally solved hank model with aggregate risk. Technical report, University of Copenhagen.

Druedahl, J. and Jørgensen, T. H. (2017). A general endogenous grid method for multi-dimensional models with non-convexities and constraints. *Journal of Economic Dynamics and Control*, 74:87–107.

Druedahl, J. and Ropke, J. (2026). Deep learning algorithms for solving convex finite-horizon models. Technical report, University of Copenhagen.

Duarte, V. (2018). Machine learning for continuous-time economics. *Available at SSRN 3012602*.

Duarte, V., Fonseca, J., Goodman, A. S., and Parker, J. A. (2021). Simple allocation rules and optimal portfolio choice over the lifecycle. Working Paper 29559, National Bureau of Economic Research.

Duffy, J. and McNelis, P. D. (2001). Approximating and simulating the stochastic growth model: Parameterized expectations, neural networks, and the genetic algorithm. *Journal of Economic Dynamics and Control*, 25(9):1273–1303.

Fella, G. (2014). A generalized endogenous grid method for non-smooth and non-concave problems. *Review of Economic Dynamics*, 17(2):329–344.

Fernández-Villaverde, J., Hurtado, S., and Nuno, G. (2023). Financial frictions and the wealth distribution. *Econometrica*, 91(3):869–901.

Folini, D., Friedl, A., Kübler, F., and Scheidegger, S. (2025). The climate in climate economics. *Review of Economic Studies*, 92(1):299–338.

Friedl, A., Kübler, F., Scheidegger, S., and Usui, T. (2023). Deep uncertainty quantification: With an application to integrated assessment models. Technical report, Working paper, University of Lausanne.

Gopalakrishna, G. (2021). Aliens and continuous time economies. *Swiss Finance Institute Research Paper*, (21-34).

Gopalakrishna, G., Gu, Z., and Payne, J. (2024). Asset pricing, participation constraints, and inequality. Technical report, Princeton Working Paper.

Grohs, P., Hornung, F., Jentzen, A., and Von Wurstemberger, P. (2018). A proof that artificial neural networks overcome the curse of dimensionality in the numerical approximation of black-scholes partial differential equations. *arXiv preprint arXiv:1809.02362*.

Gu, Z., Lauriere, M., Merkel, S., and Payne, J. (2023). Deep learning solutions to master equations for continuous time heterogeneous agent macroeconomic models.

Guarda, S. (2025). Narrow and short beliefs in macroeconomics with heterogeneous agents.

Han, J., Yang, Y., et al. (2022). Deepham: A global solution method for heterogeneous agent models with aggregate shocks. *arXiv preprint arXiv:2112.14377*.

Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of economic Dynamics and Control*, 17(5-6):953–969.

Imrohoroğlu, A. (1989). Cost of business cycles with indivisibilities and liquidity constraints. *Journal of Political economy*, 97(6):1364–1383.

Iskhakov, F., Jørgensen, T. H., Rust, J., and Schjerning, B. (2017). The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks. *Quantitative Economics*, 8(2):317–365.

Jentzen, A., Salimova, D., and Welti, T. (2018). A proof that deep artificial neural networks overcome the curse of dimensionality in the numerical approximation of kolmogorov partial differential equations with constant diffusion and nonlinear drift coefficients. *arXiv preprint arXiv:1809.07321*.

Jiang, H. (1999). Global convergence analysis of the generalized newton and gauss-newton methods of the fischer-burmeister equation for the complementarity problem. *Mathematics of Operations Research*, 24(3):529–543.

Judd, K. L. (1998). *Numerical methods in economics*. MIT press.

Judd, K. L. and Solnick, A. (1994). Numerical dynamic programming with shape-preserving splines. *Report.[657]*.

Jungerman, W. (2023). Dynamic monopsony and human capital. Technical report, mimeo.

Kahou, M. E., Fernández-Villaverde, J., Gómez-Cardona, S., Perla, J., and Rosa, J. (2022). Spooky boundaries at a distance: Exploring transversality and stability with deep learning.

Kahou, M. E., Fernández-Villaverde, J., Perla, J., and Sood, A. (2021). Exploiting symmetry in high-dimensional dynamic programming. Working Paper 28981, National Bureau of Economic Research.

Kase, H., Melosi, L., and Rottner, M. (2023). Estimating nonlinear heterogeneous agents models with neural networks. *CEPR Discussion Paper No. DP17391*.

Kase, H., Rottner, M., and Stohler, F. (2025). Generative economic modeling. Technical report, Bank for International Settlements.

Khan, A. and Thomas, J. K. (2008). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica*, 76(2):395–436.

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Krueger, D. and Kubler, F. (2004). Computing equilibrium in olg models with stochastic production. *Journal of Economic Dynamics and Control*, 28(7):1411–1436.

Krusell, P. and Smith, Jr, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy*, 106(5):867–896.

Lee, H. (2025). Global nonlinear solutions in sequence space and the generalized transition function. *Working paper*.

Levintal, O. (2018). Taylor projection: A new solution method for dynamic general equilibrium models. *International Economic Review*, 59(3):1345–1373.

Maliar, L., Maliar, S., and Winant, P. (2021). Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122:76–101.

Moll, B. (2025). The trouble with rational expectations in heterogeneous agent models: A challenge for macroeconomics. *The Economic Journal*, page ueaf104.

Norets, A. (2012). Estimation of dynamic discrete choice models using artificial neural network approximations. *Econometric Reviews*, 31(1):84–106.

Payne, J., Rebei, A., and Yang, Y. (2024). Deep learning for search and matching models. *Available at SSRN*.

Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, 33(3):649–665.

Reiter, M. (2025). Neural networks vs. traditional methods: A hybrid nonlinear approach for solving large olg models 1. *Available at SSRN 5203929*.

Rouwenhorst, K. G. (1995). Asset pricing implications of equilibrium business cycle models. *Frontiers of business cycle research*, 1:294–330.

Sauzet, M. (2021). Projection methods via neural networks for continuous-time models. *Available at SSRN 3981838*.

Sun, J. E. (2025a). Continuation value is all you need: "drop-in" deep learning for heterogeneous-agent models with aggregate shocks.

Sun, J. E. (2025b). The distributional consequences of climate change: Housing wealth, expectations, and uncertainty.

Tauchen, G. (1986). Statistical properties of generalized method-of-moments estimators of structural parameters obtained from financial market data. *Journal of Business & Economic Statistics*, 4(4):397–416.

Valaitis, V. and Villa, A. T. (2024). A machine learning projection method for macro-finance models. *Quantitative Economics*, 15(1):145–173.

Yang, Y., Wang, C., Schaab, A., and Moll, B. (2025). Structural reinforcement learning for heterogeneous agent macroeconomics. *arXiv preprint arXiv:2512.18892*.

Yokomoto, Y. (2025). Accelerating the krusell-smith algorithm with deep learning.

Young, E. R. (2010). Solving the incomplete markets model with aggregate uncertainty using the krusell–smith algorithm and non-stochastic simulations. *Journal of Economic Dynamics and Control*, 34(1):36–41.

Zhong, Y. (2023). Operator learning in macroeconomics.

Appendix

A Additional Tables

Parameter	γ	δ	β	α	ρ^A	σ^A
Meaning	relative risk aversion	depreciation of capital	patience	capital share in production	persistence of log TFP	std dev of innovations log TFP
Value	2	0.1	0.95	$\frac{1}{3}$	0.8	0.03

Table A1: Parameter values chosen for the Brock and Mirman (1972) model.

Parameter	N^{quad}	T	$N^{\text{hidden 1}}$	$N^{\text{hidden 2}}$	$N^{\text{hidden 3}}$	N^{output}
Meaning	Quad. nodes	Length of hist. of shocks	# nodes layer 1 (activation)	# nodes layer 2 (activation)	# nodes layer 3 (activation)	# nodes output layer (activation)
Value	8	100	128 (gelu)	128 (gelu)	128 (gelu)	1 (sigmoid)

Parameter	N^{data}	N^{mb}	N^{episodes}	Optimizer	α^{learn}
Meaning	States per episode	mini-batch size	# episodes	Optimizer	learning rate
Value	4096	256	40,000	Adam	10^{-5}

Table A2: Hyperparameter values chosen for the neural network to solve the Brock and Mirman (1972) model.

Parameter		Value
H	Number of cohorts	72
γ	Relative risk aversion	2
β	Patience	0.96
B	Bond supply	0.75
$\underline{b}$	Borrowing constr. bond	0
$\underline{k}$	Borrowing constr. capital	0
ξ^{adj}	Adj. cost on capital	0.1
ρ	Persistence of log TFP	0.85
σ^A	Std. dev. innovations to log TFP	0.03
δ	Depreciation of capital in normal times	0.1
ρ^δ	Persistence of depr. in disaster	0
σ^δ	Std. dev. of innovations to depreciation	0.2
μ^δ	Mean depreciation during disasters	-1.10
$\pi^{n\to n}$	Prob. to remain in normal times	0.94
$\pi^{d\to d}$	Prob. to remain in the disaster regime	2/3

Table A3: Parameter values for the OLG model in section 4.

Parameter		Value
N^{quad}	Number quadrature nodes TFP / depreciation	4 / 4
T	Length of history of shocks	144
$N^{\text{hidden 1}}$	# neurons in the first hidden layer (activation)	720 (gelu)
$N^{\text{hidden 2}}$	# neurons in the second hidden layer (activation)	720 (gelu)
$N^{\text{hidden 3}}$	# neurons in the third hidden layer (activation)	720 (gelu)
N^{output}	# neurons in the output layer (activation)	214 (None)
N^{data}	States per episode	8192
N^{mb}	mini-batch size	64
N^{episodes}, step 1	Training episodes (capital only)	512
N^{episodes}, step 2	Training episodes (pretrain bond price)	1536
$N^{\text{bond steps}}$, step 3	Number of incremental increases in bond supply	4
N^{episodes}, step 3	Training episodes (for each intermediate bond supply)	16384
N^{episodes}, step 4	Training episodes (final model)	32768
Optimizer	Optimizer	Adam
α^{learn}	Learning rate	10^{-5}

Table A4: Hyper-parameter values for the OLG model in section 4.

Parameter		Value
γ	Relative risk aversion	2
β	Patience	0.95
$\underline{\theta}$	Borrowing constraint	0
ρ^e	Persistence id. income process	0.871
σ^e	Std. dev. id. income process	0.246
ρ^A	Persistence of aggregate TFP	0.8145
σ_L^A, σ_H^A	Std. dev. innovations to TFP	1.24%, 1.99%
δ	Depreciation of capital	0.1
α	Capital share in production	0.25
ζ	Returns to scale	0.25
Γ	Survival rate of firms	0.965
z	Idiosyncratic firm productivity	[0.5, 1.0, 1.5]
Π_z^L, Π_z^H	Transition matrices for id. firm prod.	See text
$\Pi_{L,L}^U$, $\Pi_{H,H}^U$	Persistence of uncertainty regimes	0.90, 0.79
ξ^{up}, ξ^{down}, s	Adjustment costs firms	1.0, 2.5, 400

Table A5: Parameter values for the economy with heterogeneous firms and heterogeneous households.

Parameter		Value
T	Length of history of shocks (per shock)	300
N^{input}	# nodes input layer	600
$N^{\text{hidden 1}}$	# neurons in the first hidden layer (activation)	1024 (gelu)
$N^{\text{hidden 2}}$	# neurons in the second hidden layer (activation)	1024 (gelu)
$N^{\text{hidden 3}}$	# neurons in the third hidden layer (activation)	1024 (gelu)
N^{output}	# neurons in the output layer	$\mathcal{N}^k_{\boldsymbol{\rho}}$: 600 $\mathcal{N}^\lambda_{\boldsymbol{\rho}}$: 600 $\mathcal{N}^c_{\boldsymbol{\rho}}$: 200 $\mathcal{N}^p_{\boldsymbol{\rho}}$: 1
N^k	# grid points for capital grid (log spaced)	200
N^θ	# grid points for asset grid (log spaced)	200
N^{quad}	# quadrature nodes TFP	5
N^{data}	States per episode	131072
N^{mb}	mini-batch size	128
N^{episodes}, step 1	Training episodes	100
N^{episodes}, step 2	Training episodes	100
N^{episodes}, step 3	Training episodes	100
N^{episodes}, step 4	Training episodes	500
N^{episodes}, step 5	Training episodes	16300
Optimizer	Optimizer	Adam
α^{learn}	Learning rate	10^{-6}

Table A6: Hyperparameter values for the heterogeneous firms and households model.

Parameter	Meaning	Value
H	Number of cohorts	10
γ	Relative risk aversion	2
β	Patience (per 6-year period)	$0.95^6 \approx 0.735$
α	Capital share in production	$1/3$
$\underline{k}$	Borrowing constraint	0
N_k	Asset grid points	300
$k \in [\underline{k}, \bar{k}]$	Asset grid bounds	$[0, 5]$
$\mathcal{E}^{\text{grid}}$	Idiosyncratic productivity states	$\{0.5, 1.0, 1.5\}$
Π^e	Transition matrix for e	See text
δ_n	Depreciation (normal)	$1 - (1 - 0.05)^6 \approx 0.265$
δ_d	Depreciation (disaster)	$1.5\delta_n \approx 0.397$
Π^z	Transition matrix for z	See text
ϵ	Utility cost of work at $h = 6$	1.68
σ	Taste-shock scale (logit smoothing)	0.08
y^{home}	Income when retired	0.1
ψ^b	Bequest motive strength	0.1
ν	Bequest curvature	2
$\bar{b}$	Bequest shifter	0.1

Table A7: Parameter values for the discrete-continuous choice OLG model in section 6.

B Additional Figures

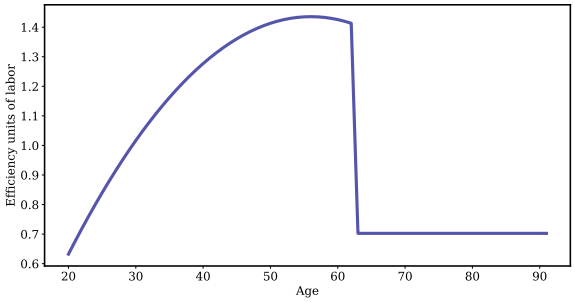


Figure A1: Efficiency units of labor over the life-cycle for the OLG model in section 4.

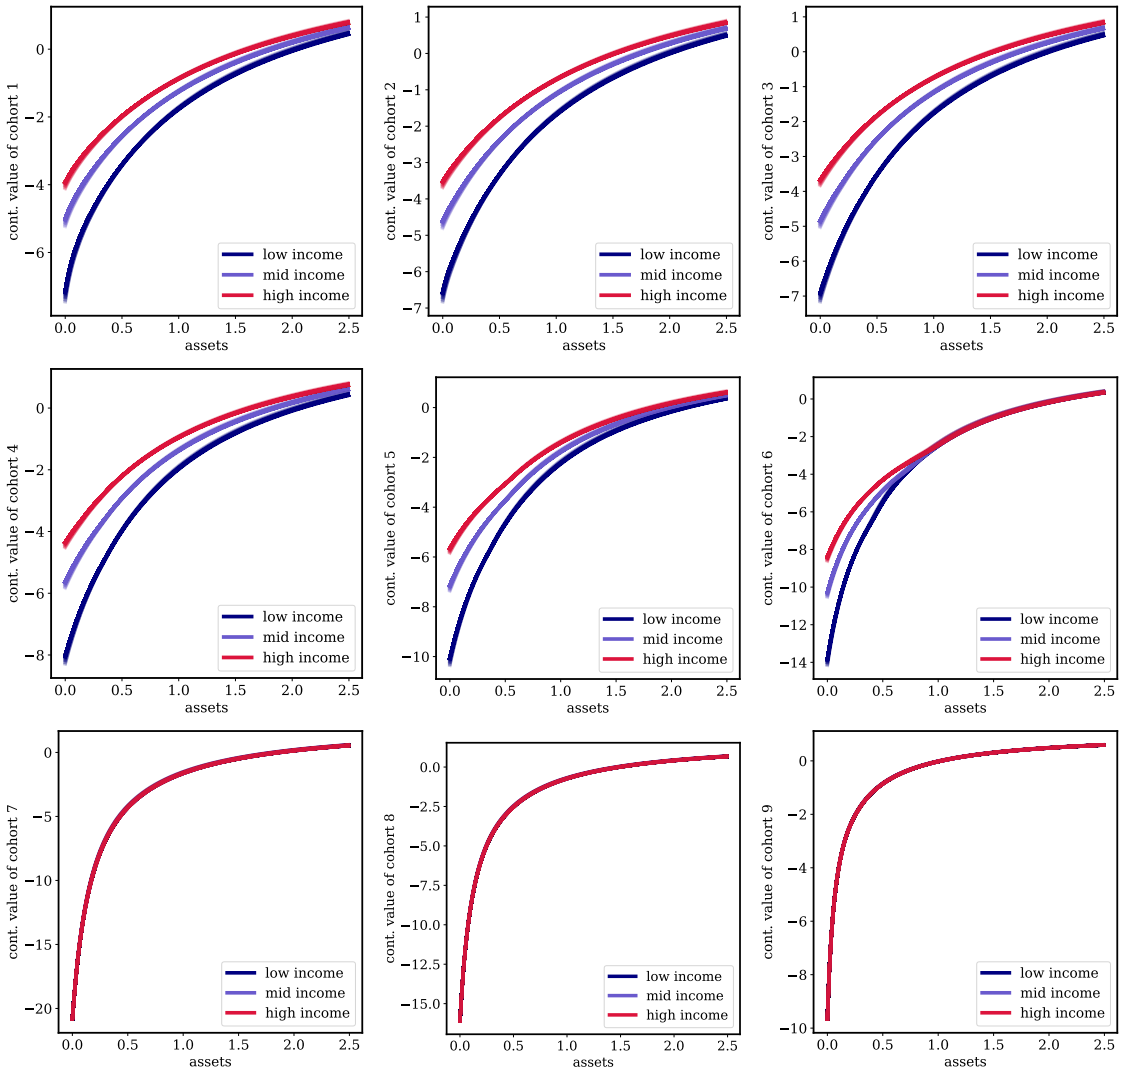


Figure A2: Continuation values by age and asset holdings for different idiosyncratic productivity and several aggregate states.

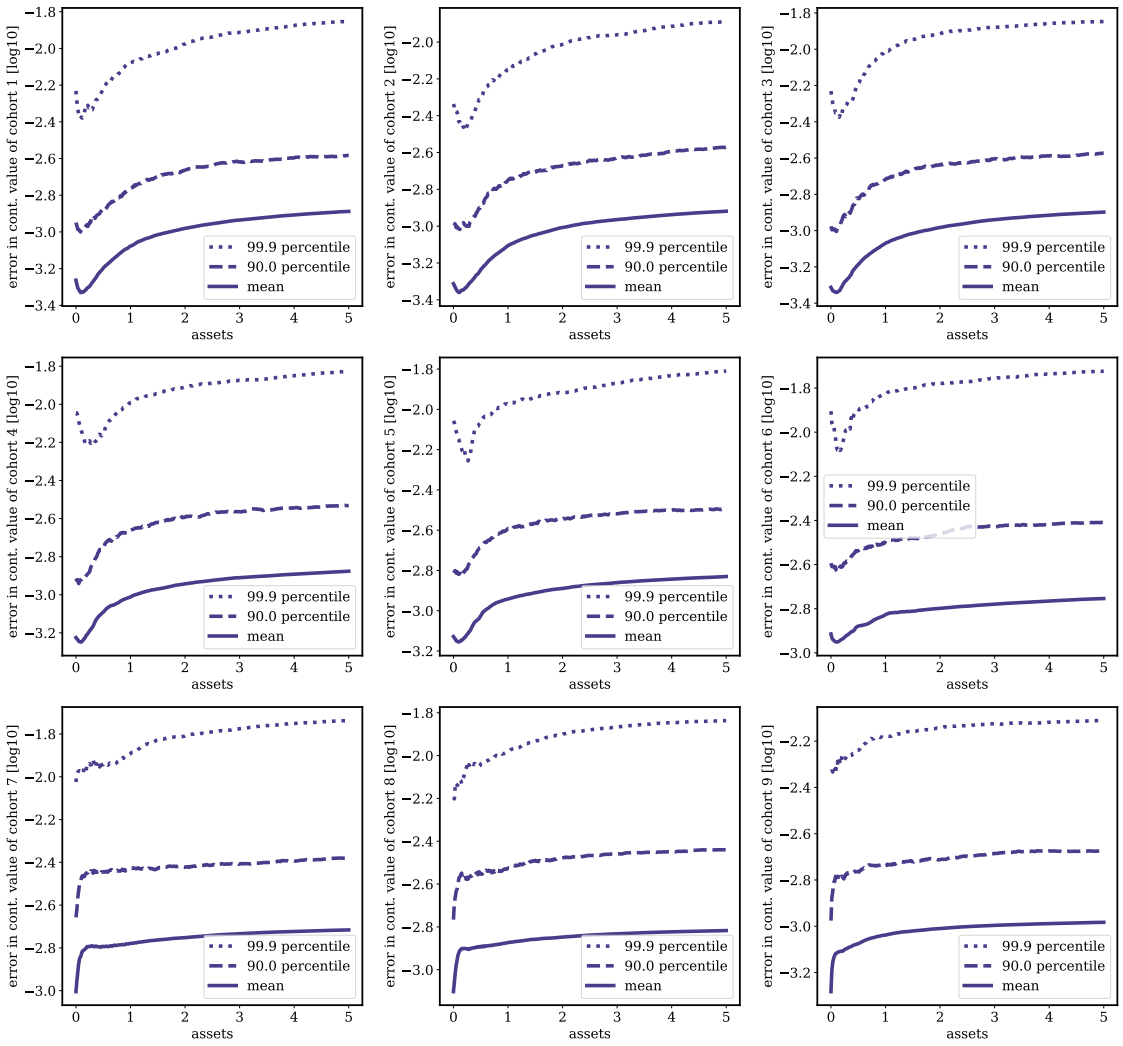


Figure A3: The distribution of (absolute) prediction errors for the continuation value by age and asset holdings.

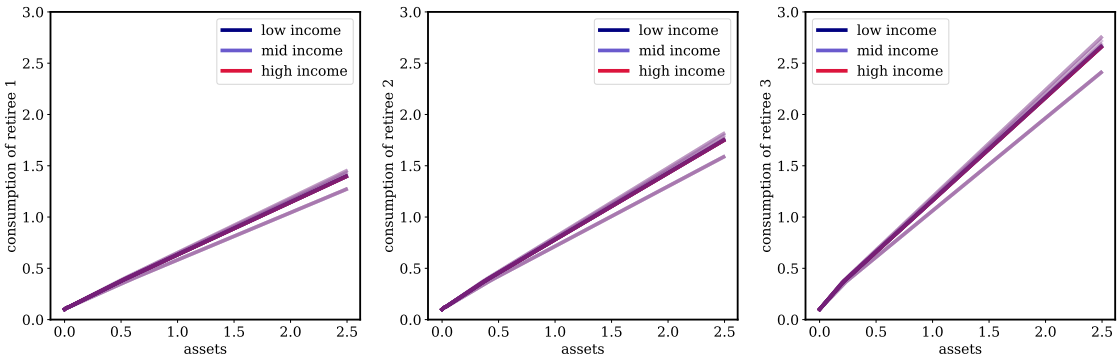


Figure A4: Consumption functions of retired cohorts. The shaded lines show the policies for different aggregate states, and the solid lines show the average taken across the sample of aggregate states.