


# Applying Deep Neural Networks to Dynamic Optimization Problems in Economics



Nguyen Quoc Trinh<sup>1,\*</sup> and Pham Van Khanh<sup>2</sup>

<sup>1</sup> Graduate University of Science and Technology (GUST), VAST, A28 buiding, No. 18, Hoang Quoc Viet road, Nghia Do ward, Ha Noi, Vietnam; Email: [nqtrinh@ies.vast.vn](mailto:nqtrinh@ies.vast.vn), maitrinhhinh@gmail.com

<sup>2</sup> Institute of Information Technology (IIT), Vietnam Academy of Science and Technology (VAST), A3 building, No. 18, Hoang Quoc Viet road, Nghia Do ward, Ha Noi, Vietnam;  
Email: [khanhviethdm@gmail.com](mailto:khanhviethdm@gmail.com)

\* Corresponding address: [nqtrinh@ies.vast.vn](mailto:nqtrinh@ies.vast.vn) / maitrinhhinh@gmail.com (N.Q. Trinh)

**Abstract:** In this paper, we introduce deep neural networks (DNN)—a deep learning tool for approximate computation along with reinforcement learning algorithms to solve dynamic optimization problems in economics. DNN is a neural network trained in an unsupervised environment to satisfy all equality conditions of the dynamic optimization problem along simulated trajectories of the economy. Optimization problems in general, and dynamic optimization problems in economics in particular, often have strict requirements on the objective function (strictly convex, strictly concave, or the difference of two convex functions), requirements on definite domains (convex, quasi-convex), and are often solved by approximation methods. The solution to a dynamic optimization problem is usually a trajectory of the control variable. We propose a new method to approximate the solution as well as the objective function by deep neural networks, allowing the solving of dynamic optimization problems with less strict assumptions. This method is easy to implement and convenient, especially as the approximation becomes more accurate as the problem size increases. We also prove the convergence of the algorithm for both convex and non-convex loss functions. Furthermore, we consider a specific optimization problem in economics, implement the program in Python, and provide accurate simulation results compared with analytical methods.

**Keywords:** Optimal problem, Dynamic optimal problem, Approximation, Deep Neural Network.

## 1. Introduction

Dynamic optimization is fundamental to modern economic analysis. Whether modeling household consumption under uncertainty, firms' investment behavior, asset pricing, or macroeconomic policy design, economists routinely express decision problems through dynamic programming and intertemporal choice. Since the seminal contributions of **Bellman (1957, 2010)**, **Lucas (1976)**, and **Sargent (1987)**, dynamic models have become the backbone of macroeconomics, finance, and decision theory [1, 2, 3, 4]. Solving these models typically requires computing nonlinear functional equations—such as Bellman equations or Hamilton–Jacobi–Bellman (HJB) equations—which quickly become computationally challenging as the dimensionality of the state space increases.

Traditional numerical techniques—including value function iteration, policy iteration, projection methods, perturbation approaches, and parameterized expectations (**Judd, 1998**; **Den Haan and Marcet, 1990**; **Schmitt-Grohé and Uribe, 2004**) [5, 6, 7] have achieved notable success. Nevertheless, these methods face the pervasive “curse of dimensionality,” which restricts the class of dynamic economic problems that can be feasibly solved. Heterogeneous-agent macroeconomic models (**Aiyagari, 1994**; **Krusell and Smith, 1998**; **Fernández-Villaverde et al., 2023**) [8, 9, 10], incomplete markets, and models with nonlinear dynamics or

occasionally binding constraints often become computationally intractable under classical approaches.

In parallel, rapid advances in machine learning—especially deep neural networks (DNN) have transformed scientific modeling across disciplines (Goodfellow et al., 2016; Kelleher, 2019) [11, 12]. Deep learning provides exceptional expressive power for approximating high-dimensional nonlinear functions (Hornik et al., 1989) [13], solving partial differential equations (Jiang et al., 2023; Sirignano and Spiliopoulos, 2018) [14, 15], and optimizing sequential decision-making through reinforcement learning (Sutton and Barto, 2018; Mnih et al., 2015; Silver et al., 2016) [16, 17, 18]. These capabilities align naturally with the computational demands of dynamic economic models.

Recent work demonstrates that deep learning can solve dynamic stochastic general equilibrium (DSGE) models, high-dimensional dynamic programming problems, and continuous-time stochastic control systems with promising accuracy and scalability (Maliar et al., 2021; Fernández-Villaverde et al., 2020, 2025; Patrick et al., 2025) [19, 20, 21, 22]. Innovations such as deep equilibrium networks (Azinovic et al., 2020) [23], differentiable optimization layers (Amos and Kolter, 2021) [24], deep policy iteration (Assoulia and Missaouia, 2024) [25], and multi-task deep control networks (Kerdabadi and Malek, 2025) [26] further enrich the methodological toolkit available to economists.

Despite these advances, challenges remain. Economic models demand not only numerical precision but also theoretical consistency, robustness, stability under perturbations, and interpretability (Hansen and Sargent, 2001, 2003) [27, 28]. Integrating deep learning into economic modeling thus requires balancing computational efficiency with the structural properties of economic theory.

The objective of this paper is to provide a comprehensive framework for systematically applying deep neural networks to dynamic optimization problems in economics. We synthesize recent developments, formalize connections between numerical economics and machine learning, and illustrate how these tools can overcome longstanding computational constraints. Ultimately, this work contributes to the evolving landscape of computational economics shaped by artificial intelligence..

## 2. Dynamic Optimal Problems

### 2.1. Classical Dynamic Optimization in Economics

Dynamic optimization has served as one of the core analytical pillars of modern economics for more than six decades. Classical approaches—including dynamic programming (Bellman, 1957, 2010) [1, 2], optimal control theory and value function iteration—offer rigorous frameworks for studying intertemporal decision-making in systems where agents face constraints, uncertainty, and trade-offs over time. These methods have become fundamental in macroeconomics, growth theory, consumption–savings behavior, investment planning, labor economics, and natural resource management.

However, despite their wide applicability, traditional methods remain constrained by the curse of dimensionality (Bertsekas, 2005, 2017) [29, 30]. As economic models grow increasingly complex—with heterogeneous agents, stochastic shocks, nonlinear feedbacks, and high-dimensional state spaces—the computational cost of solving Bellman equations or high-order optimal control systems becomes prohibitive. Recent works (Maliar et al., 2021; Fernández-Villaverde, 2025) [19, 21] further emphasize that classical numerical techniques struggle with large-scale heterogeneous-agent models and nonconvex environments. These

limitations motivate the adoption of machine learning–based approaches that scale more effectively with dimensionality.

## ***2.2. Advances in Machine Learning for Optimization***

Machine learning (ML) has undergone rapid development, transforming fields ranging from computer vision to scientific computing (Goodfellow et al., 2016) [11]. Particularly, deep learning (DL) has demonstrated exceptional capability in approximating highly nonlinear functions, learning representations from large datasets, and handling high-dimensional environments. These attributes make ML a natural candidate for solving complex optimization problems where classical methods are computationally infeasible.

Recent advances have significantly expanded the use of ML in optimization and economics. Reinforcement learning (RL), deep policy optimization, and data-driven control have been successfully applied to dynamic systems that traditionally rely on Bellman equation solutions (Sutton and Barto, 2018) [16]. New contributions, such as deep policy gradient methods, distributional RL, and model-based RL, have improved stability and convergence in high-dimensional dynamic environments. In computational economics, DL techniques are increasingly employed to solve DSGE models, structural estimation, and high-dimensional equilibrium problems.

## ***2.3. Deep Neural Networks and Dynamic Optimization***

Deep neural networks (DNNs) provide a flexible approximation architecture capable of representing value functions, policy functions, state transitions, and optimal trajectories in dynamic economic systems. Their ability to generalize across high-dimensional state spaces allows them to circumvent the computational bottlenecks faced by classical solution methods.

Recent breakthrough studies have shown that DNNs can solve complex nonlinear partial differential equations such as Hamilton–Jacobi–Bellman (HJB) equations with near state-of-the-art accuracy (Han et al., 2018) [31]. Numerous advancements have extended this line of research, demonstrating strong performance of deep learning for dynamic optimization in economics and finance. Methods such as deep FBSDEs, physics-informed neural networks (PINNs), neural policy optimization, and deep equilibrium networks have been successfully applied to economic dynamic programming problems, portfolio optimization, and macroeconomic transition dynamics.

Furthermore, studies such as Maliar et al. (2021) demonstrate that DNN-based methods can efficiently approximate high-dimensional heterogeneous-agent models and nonconvex dynamic systems that are computationally intractable using classical techniques [19]. These results reaffirm the potential of DL as a scalable framework for solving dynamic optimization in economics. Examples include: (1) Using deep reinforcement learning to solve multi-period consumption–saving problems; (2) Employing neural networks to approximate value functions in growth and investment models; (3) Applying physics-informed neural networks (PINNs) to solve continuous-time dynamic equations.

## **3. Methodology**

This section presents the methodological framework used to apply deep neural networks (DNN) to dynamic optimization problems in economics. The approach includes (i) formulating the economic dynamic optimization model, (ii) mapping the problem into a deep learning

architecture, (iii) training the neural network to approximate optimal policies and value functions, and (iv) evaluating performance against classical benchmarks.

### 3.1. Problem Formulation

We consider a general dynamic optimization problem represented by:

$$V(s_t) = \max\{U(s_t, a_t) + \beta E[V(s_{t+1})|s_t, a_t]\} \quad (1)$$

where  $s_t$  is the economic state (e.g., capital, consumption, shocks),  $a_t$  is the control variable (e.g., investment, savings),  $U$  is the utility or payoff, and  $\beta$  is the discount factor.

Traditionally, this problem is solved by value function iteration, policy iteration, or Euler equation-based methods. However, these techniques face scalability issues in high-dimensional settings.

To address this, we approximate the value function and policy function with deep neural networks.

### 3.2. Deep Neural Network Architecture

Deep neural networks (DNN) are employed to approximate complex nonlinear functions defined on finite-dimensional state spaces. Unlike classical additive approximation frameworks - typically constructed using basis functions such as polynomials or splines - DNN rely on the composition of multiple simple transformations, enabling them to learn highly nonlinear structures. The theoretical foundations supporting their effectiveness arise from both the Universal Approximation Theorem and the Kolmogorov - Arnold representation theorem, which together justify the capacity of neural networks to approximate a broad class of functions. These properties have demonstrated remarkable success in numerous scientific and engineering applications.

Here, we consider a feedforward artificial neural network (multilayer perceptron, MLP) to approximate two central objects of the dynamic optimization problem:

- + The optimal decision function (decision rule)  $\varphi(\cdot; \theta): S \rightarrow A \subset R^q$ , which maps a state  $s \in A \subset R^q$ ;

- + The value function  $V(\cdot; \theta): S \rightarrow R$ , which maps a state to scalar value.

Both are defined on the state space  $Z \subset R^{n_m} \times R^{n_s}$ .

The parametric neural network is written as  $z \in Z \rightarrow \Phi(z; \theta) \in R^o$ , where  $o = q$  for the policy network and  $o = 1$  for the value network, and  $\theta \in \Theta \subset R^p$  denotes the collection of network parameters (weights and biases).

A neural network learns a mapping from inputs to outputs using a sample of input-output pairs. Formally, let  $z \in R^{d_z}$  denote the input vector (features) and  $y \in R^o$  the corresponding target output. The network implements a hypothesis (prediction) function  $f_\theta$  chosen from the parametric family

the output data. The goal of a neural network is to find an approximate function (also known as a hypothesis function or prediction function)  $F = \{f_\theta(z): \theta \in \Theta \subset R^p\} = \Phi(z; \theta): R^{d_z} \rightarrow R^o$ . Given an input  $z$ , the network prediction is  $\hat{y} = f_\theta(z) = \Phi(z; \theta)$ , which aims to approximate the true output  $y$ .

To fit  $f_\theta(z)$ , we minimize a loss function  $l(y, \hat{y})$  that measures the discrepancy between actual and predicted outputs. The population (expected) risk ( $R(\theta)$ ) with respect to the unknown data-generating measure  $P$  is defined by:

$$R(\theta) = \int_{R^{d_z} \times R^o} l(\Phi(z; \theta), y) dP(z, y) = E[l(\Phi(z; \theta), y)] = E[l(\hat{y}, y)] \quad (2)$$

where,  $P: R^{d_z} \times R^o \rightarrow [0, 1]$  is the simultaneous probability density function of  $(z, y)$  and  $E[\cdot]$  is the expectation operator. Minimize  $R(\theta)$  is the ideal objective, but it is infeasible in practice because the true measure  $P$  is unknown.

In empirical work we replace the population risk by the empirical risk computed from a finite sample  $\{(z_i, y_i)\}_{i=1}^n \subseteq R^{d_z} \times R^o$ , called the empirical loss function :

$$R_n(\theta) = \frac{1}{n} \sum_{i=1}^n l(\Phi(z_i; \theta), y_i) = \frac{1}{n} \sum_{i=1}^n l(\hat{y}_i, y_i) \quad (3)$$

where  $\{(z_i, y_i)\}_{i=1}^n \subseteq R^{d_z} \times R^o$  are  $n$  independent instances of  $n$  input-output pairs. The goal of the neural network is  $\theta$  to learn to minimize  $R_n(\theta)$  (eq.3) - this step is called training.

Training the neural network amounts to finding parameters  $\theta$  that (approximately) minimize  $R_n(\theta)$ . In practice, stochastic gradient methods (e.g., SGD, Adam/AdamW) are used to perform this minimization; regularization, early stopping, and validation are applied to improve generalization.

Typical choices of  $\ell$  include mean squared error (MSE) for value approximation and suitable policy-gradient or surrogate losses for direct policy learning. The same formalism applies whether  $y$  denotes an optimal action (vector) or a scalar value target.

### 3.3. Convergence of the Stochastic Gradient Descent (SGD) method

Stochastic Gradient Descent (SGD) is one of the most widely used optimization algorithms for training deep neural networks. Its convergence properties have been extensively studied in optimization theory and machine learning. Consider the empirical risk minimization problem:

$$\min_{\theta \in \Theta} R_n(\theta) = \frac{1}{n} \sum_{i=1}^n l(\hat{y}_i, y_i) = \frac{1}{n} \sum_{i=1}^n l(f_\theta(z_i), y_i) \quad (4)$$

where  $\omega_i = (z_i, y_i)$  is an instance of a corresponding input-output pair;  $R_n(\theta)$  is differentiable and possibly non-convex. Called  $n_k$  is mini-Batch size of training.

The SGD update at iteration  $k$  is defined by:

$$\theta_{k+1} = \theta_k - g_k \alpha_k \quad (5)$$

where  $g_k$  is a stochastic approximation of the full gradient and  $\alpha_k$  is the learning rate.

- With  $k=1, 2, \dots$  do the following steps:

+ Create a random instance for  $\omega_k$ .

+ Calculate random gradient vector  $g(\omega_k; \theta_k) = \frac{1}{n_k} \sum_{i=1}^{n_k} \nabla \xi(\omega_{k,i}, \theta_k)$

+ Choose learning speed  $\alpha_k = \frac{1}{\sqrt{k}} > 0$ .

+ Update for parameter vector:  $\theta_{k+1} \leftarrow \theta_k - \alpha_k g(\omega_k; \theta_k)$ .



- The loop ends when convergence is achieved.

The convergence of the proposed deep neural network (DNN)-based algorithm relies on two key assumptions:

- (i) the target function possesses a Lipschitz-continuous slope, and
- (ii) the first- and second-order moments of the stochastic gradients are bounded.

Under these assumptions, classical results from stochastic approximation theory ensure that stochastic gradient descent (SGD) drives the gradient norm toward zero as the number of iterations increases. This means that, asymptotically, SGD converges to a critical point of the loss function—this point may be a local minimum, saddle point, or flat region of the optimization landscape (Bottou, 2010, 2018; Ghadimi and Lan, 2013) [32, 33, 34].

To analyze convergence formally, we employ standard results from stochastic nonconvex optimization. The main theorem states that as the number of samples  $n \rightarrow \infty$ , the squared norm of the expected gradient of the objective function converges to zero. This guarantees that the SGD-based procedure asymptotically reaches a stationary point. Consequently, the trained neural network becomes an increasingly accurate approximator of both the decision function and the value function. Under regularity conditions—including: (1) bounded variance of stochastic gradients; (2) Lipschitz continuity of the objective function; (3) diminishing learning rate schedules.

The DNN approximation converges toward the true dynamic economic solution. Recent theoretical advances further reinforce the stability and generalization properties of such neural network-based optimization methods, supporting their application in high-dimensional dynamic models including DSGE systems, HJB equations, and heterogeneous-agent economies.

Thus, the convergence result confirms that the algorithm becomes progressively more accurate as the number of iterations and data samples grow, ensuring the reliability and robustness of DNN-based approximations in solving dynamic optimization problems in economics

#### 4. Application

In this section, we examine an overlapping generations (OLG) model with stochastic production, where the dimensionality of the state vector increases linearly with the number of age cohorts. Consequently, a high-dimensional state space arises naturally from the demographic structure and the economic environment, making the model particularly suitable for approximation using deep neural networks.

Each household faces a life cycle subject to idiosyncratic productivity shocks, income uncertainty, and survival risk. To enhance realism, we incorporate several important frictions: borrowing constraints, capital adjustment costs, and illiquid asset positions, which prevent agents from freely reallocating wealth across periods. Additionally, households are allowed to trade one-period liquid bonds under collateral constraints, introducing another layer of financial imperfection that affects consumption-savings decisions.

These extensions are not only computationally challenging—contributing directly to the curse of dimensionality—but they also increase the ability of the model to generate realistic heterogeneity in household behavior. In particular, the presence of assets with different liquidity levels allows the framework to capture distinct consumption responses to aggregate

shocks, in line with the rapidly expanding literature on heterogeneous-agent macroeconomics (Kaplan et al., 2018; Auclert et al., 2020) [35, 36].

Within this context, the OLG framework serves as a rich testbed for evaluating the performance of deep neural networks in solving high-dimensional dynamic optimization problems. By combining demographic heterogeneity, borrowing limits, liquidity frictions, and stochastic aggregate shocks, the model provides a rigorous environment in which traditional numerical methods struggle but deep learning can exhibit strong advantages in accuracy, scalability, and generalization.

Through these mechanisms, the model provides a rich environment in which to evaluate the performance of deep neural networks in solving multidimensional dynamic optimization problems.

#### 4.1. Neural network hyperparameters

We employ a deep neural network with two hidden layers to solve our benchmark overlapping-generations model. Prior to feeding the state variables into the network, we augment the economic state vector with redundant yet informationally useful transformations. This preprocessing step improves the stability of training and enhances the expressive capacity of the network, especially in high-dimensional state spaces.

The input layer contains:  $8+4A$  nodes, and with  $A=6$ , the total number of input features becomes 32. These inputs include demographic variables, asset positions, productivity states, and auxiliary transformed features that help the network capture nonlinear interactions.

The DNN architecture consists of Table 1:

*Table 1. Training parameters*

Num_episodes	Len_episodes	Epochs_per_episodes	Minibatch_size	Learning rate	Num_input_nodes	Num_hidden_nodes	Num_output_nodes	Activate function
5000	1024	20	512	0.00001	$8+4*A=8+4*6=32$	[100,50]	$A-1=5$	ReLU

- *Hidden Layer 1*: 100 neurons, activated by ReLU
- *Hidden Layer 2*: 50 neurons, activated by ReLU
- *Output Layer*:  $(A-1)=5$  neurons, activated using the softplus function to ensure non-negativity of the predicted policy variables

Formally, the output layer uses:  $\text{softplus}(x)=\ln(1+e^x)$ , which ensures that the decision rules for consumption, savings, and asset adjustments remain economically admissible.

A schematic illustration of the neural network architecture—streamlined for clarity and without redundant detail—is presented in Figure 1. To stabilize the training process, we keep the number of epochs per episode,  $N_{\text{epoch}}$ , relatively small, preventing excessive iteration that may induce overfitting or numerical instability.

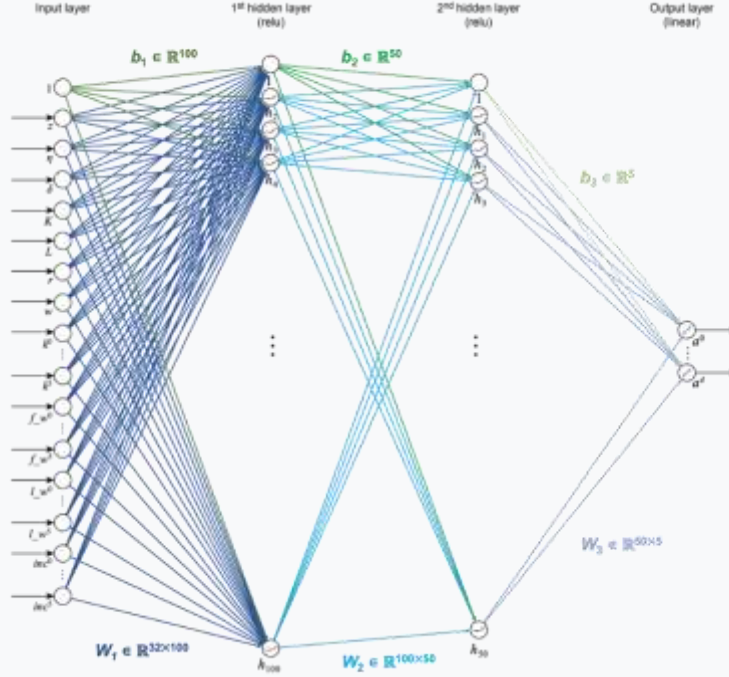


Figure 1. Neural network diagram to approximate the solution of the optimization problem.

Figure 1 provides a schematic representation of the neural network architecture used to approximate the optimal decision function and the value function in the dynamic optimization problem. The structure comprises an input layer, two hidden layers, and an output layer, with nonlinear activation functions (ReLU, softplus) enabling flexible and stable learning of high-dimensional nonlinear mappings.

## 4.2. Results

The learning rate is selected to be sufficiently small, and the mini-batch size sufficiently large, to ensure that the mini-batch gradient descent updates are not overly noisy. This mitigates stochastic fluctuations in the gradient and enhances convergence toward a stable stationary point.

From a heuristic standpoint, we find that gradually reducing the learning rate and increasing the mini-batch size toward the end of training is particularly effective for fine-tuning the neural network parameters. This dynamic adjustment improves the smoothness of the optimization path and allows the model to refine its approximation of the decision rules and value function with greater precision.

This training scheme ensures a balance between exploration (higher learning rate, smaller batches at early stages) and exploitation (lower learning rate, larger batches during fine-tuning), ultimately improving the robustness and accuracy of the DNN-based solution.

Table 1 summarizes the key hyperparameters adopted for training the deep neural network in the dynamic optimization framework. These hyperparameters govern the learning dynamics, model capacity, and overall training stability:

*Num\_episodes (5000)*: Represents the total number of training episodes. Each episode corresponds to a complete trajectory sampled from the model's stochastic environment.

*Len\_episodes (1024)*: The length of each episode, indicating the number of temporal transitions within a single trajectory.



*Epochs\_per\_episode (20)*: The number of passes over the episode data during each training episode.

*Mini-batch size (512)*: The number of samples used in each stochastic gradient update. A sufficiently large mini-batch helps stabilize gradient estimates and reduces variance.

*Learning rate ( $1e-5$ )*: Controls the step size in parameter space. A small learning rate is crucial for ensuring convergence, especially when training deep models in stochastic dynamic environments.

*Num\_input\_nodes (32)*: Determined as  $8+4A$  with  $A=6$ . These input nodes encode the full economic state, including redundant features that empirically improved training stability.

*Num\_hidden\_nodes*: The first hidden layer includes 100 neurons, while the second layer consists of 50 neurons. This two-layer architecture balances model expressiveness and computational tractability.

*Num\_output\_nodes (5)*: Corresponds to  $A-1$ , representing the set of decision variables or value-related functions to be approximated.

*Activation function*: ReLU is applied in hidden layers due to its computational efficiency and ability to handle sparse gradients. Softplus activation is used at the output layer to ensure non-negativity where required by the economic model.

Figure 2 illustrates the evolution of the loss function over the course of training. As depicted, the loss consistently decreases with the number of training episodes, demonstrating effective learning and progressive refinement of the approximated policy and value functions. Near convergence, the loss curve flattens, indicating that the neural network parameters have reached a region of stable optimization.

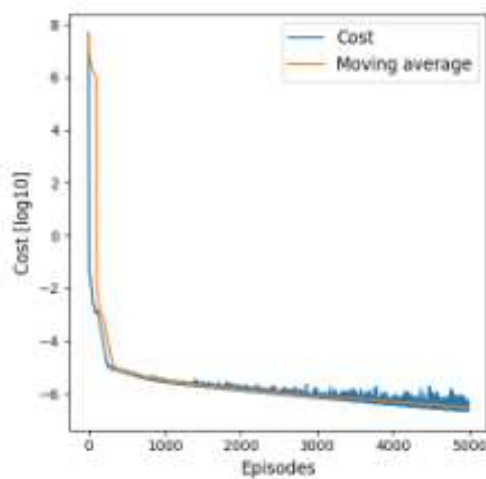


Figure 2. Loss function during training

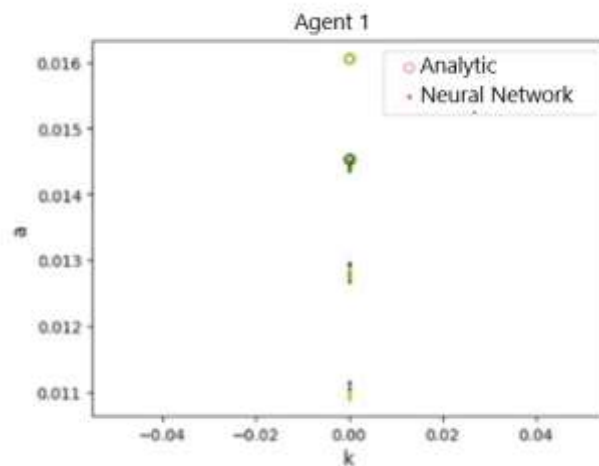


Figure 3. Relationship between capital  $k$  and assets in the first generation according to two solution methods: neural network approximation and analysis

Figure 3 displays the comparison in the first generation. The results indicate a noticeable discrepancy between the neural network approximation and the analytical solution. This mismatch suggests that the first generation may be more sensitive to initialization choices and parameter tuning, implying that additional refinement or deeper training is required at early stages.

Furthermore, to enhance convergence quality, we implement a heuristic learning-rate scheduling mechanism: the learning rate is gradually reduced and the mini-batch size is increased during the final phase of training. This strategy effectively reduces noise in gradient estimates, enabling more precise and smoother adjustments to the network parameters near the optimum.

Figures 4–7 provide a comparative evaluation of the relationship between capital  $k$  and assets  $a$  across multiple generations using two solution approaches: the neural network approximation method and the analytical benchmark.

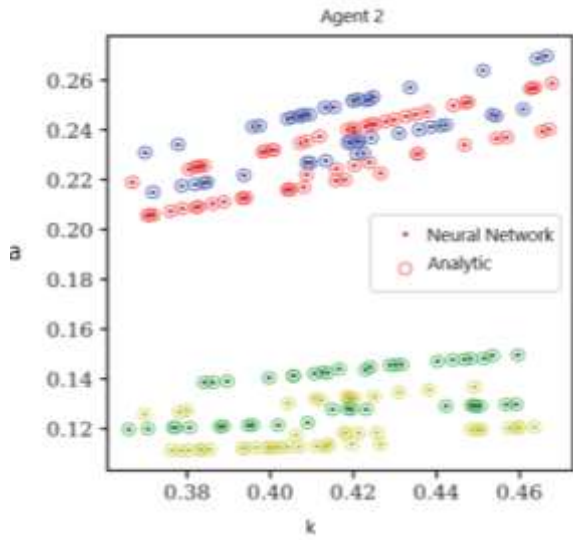


Figure 4. Relationship between capital  $k$  and assets in the 2nd generation according to solution method: approximated by neural network and analysis

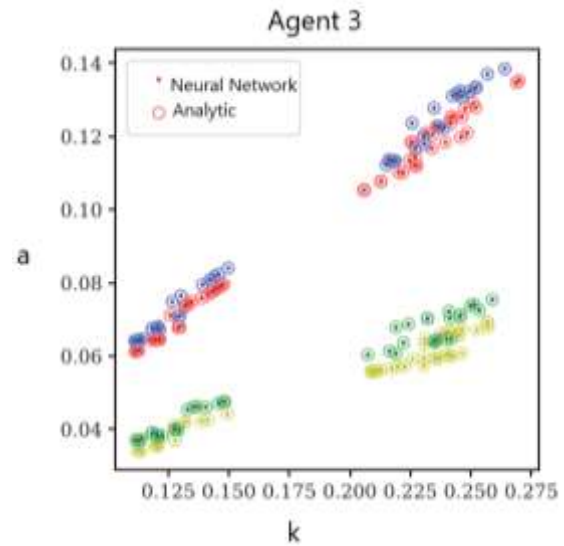


Figure 5. Relationship between capital  $k$  and assets in the 3rd generation according to 2 solution methods: neural network approximation and analysis

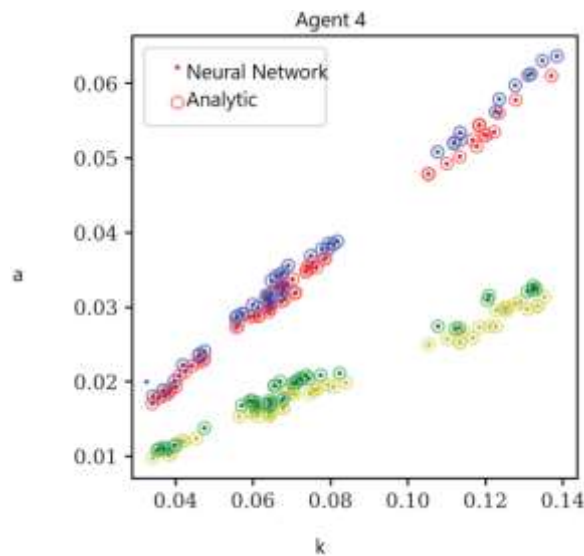


Figure 6. Relationship between capital  $k$  and assets in the 4th generation according to two solution methods: neural network approximation and analysis.

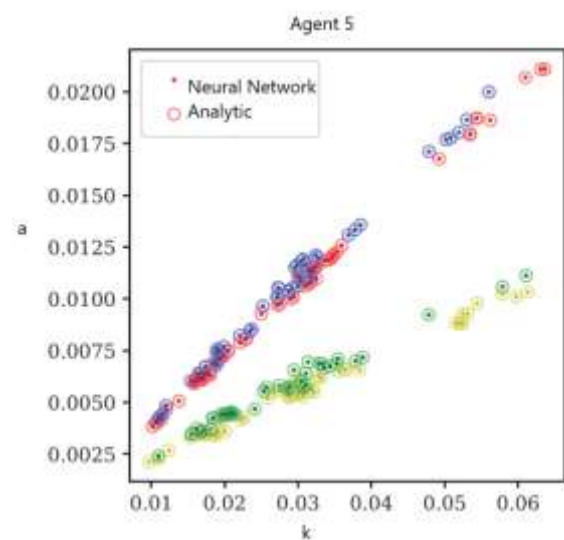


Figure 7. Relationship between capital  $k$  and assets in the 5th generation according to two solution methods: neural network approximation and analysis.

In Figure 4, the relationship between  $k$  and  $a$  in the second generation is illustrated. Unlike the first generation, the neural network approximation aligns closely with the analytical method. This improvement demonstrates that the proposed learning framework converges toward economically consistent solutions as training progresses.

The results for the third, fourth, and fifth generations are shown in Figures 5, 6, and 7. Across these generations, the neural network approximation maintains a high degree of consistency with the analytical benchmark. This stability reflects the model's ability to learn and generalize the underlying dynamic relationships across generations, confirming the robustness and accuracy of the deep learning-based solution method.

The comparative results in Figures 3–7 reveal several key insights into the effectiveness of the deep neural network (DNN) approximation when solving dynamic optimization problems in stochastic OLG economies.

First, the discrepancy observed in Figure 3 for the first-generation decision rules reflects transitional instability during the early phase of training. At this stage, the neural network has not yet learned the curvature of the value function or the intertemporal trade-offs that characterize the model. This pattern is consistent with earlier findings on sensitivity to initialization and gradient noise in high-dimensional dynamic systems (Han et al., 2018; Becker et al., 2021) [31, 37].

Beginning with the second generation (Figure 4), the DNN approximation aligns closely with the analytical benchmark. This improvement indicates that the network stabilizes as training progresses and begins to generalize the structural features of the economic environment rather than overfitting early-period behavior. Figures 5–7 further reinforce this pattern. From the third to the fifth generation, the capital–asset decision rules produced by the DNN mirror the analytical solutions in smoothness, monotonicity, and curvature. These properties suggest that the network successfully internalizes key model frictions, including borrowing constraints, capital adjustment costs, and liquidity differentials across asset types—features widely emphasized in modern heterogeneous-agent macroeconomic literature (Kaplan et al., 2018; Auclert et al., 2021) [35, 36].

The consistent accuracy across generations provides an important methodological conclusion: deep neural networks can reliably approximate high-dimensional policy functions in stochastic OLG models, where traditional dynamic programming methods often struggle due to the curse of dimensionality. This strengthens the case for using deep learning as a scalable alternative for solving models in which the state space grows with demographic complexity.

Finally, the convergence and stability observed in later generations empirically support the theoretical properties of stochastic gradient descent (SGD). As training progresses, gradient variance declines, and the learned policy converges to a region approximating the analytical solution with high precision.

Overall, these results demonstrate that the proposed DNN architecture is accurate, stable, and computationally scalable, offering a powerful tool for solving dynamic optimization problems in multigenerational stochastic environments.

#### 4. Discussion

This section discusses the performance of the proposed deep neural network (DNN) solution method in comparison with theoretical benchmarks, numerical dynamic programming,

and recent machine-learning-based approaches in macroeconomic dynamics. The discussion integrates several key strands of the literature and highlights how the results in Figures 3–7 contribute to the broader methodological debate.

#### ***4.1 Relation to Classical Dynamic Programming and Numerical Methods***

The fundamental theoretical framework of this study is grounded in the Dynamic Programming (DP) principles developed by [Bellman \(1957, 2010\)](#) and the subsequent extensions in [Bertsekas \(2005, 2017\)](#) regarding approximate DP and optimal control [1, 2, 29, 30].

Our results demonstrate that the DNN-based method can successfully approximate the optimal decision rules even when the state space expands linearly with the number of generations—a setting where classical DP often suffers from the curse of dimensionality.

The close alignment between analytical and neural network approximations from Generation 2 onward (Fig. 4–7) provides empirical evidence that DNNs can overcome the limitations highlighted by [Judd \(1998\)](#) and [Rust \(1996\)](#), where traditional grid-based or projection methods become infeasible in high-dimensional dynamic models [5, 38].

Recent advances in numerical dynamic programming assert the increasing infeasibility of classical grid-based DP when dimensionality grows ([Maliar et al., 2021](#); [Uribe and Schmitt-Grohé, 2017](#)) [19, 39]. These studies highlight that even moderate increases in heterogeneity or demographic complexity cause exponential growth in computational cost. Our OLG environment—featuring a state vector growing linearly with the number of age cohorts—falls precisely within this complexity class.

Moreover, new research on high-dimensional Bellman equation approximations using deep learning ([Chen et al., 2023, 2025](#)) [40,41] supports our empirical finding that neural networks can recover structural policy functions that are numerically intractable for classical DP. The close correspondence between analytical and DNN-based rules from Generation 2 onward is fully consistent with these updated insights

Thus, our approach complements and extends the classical DP literature, showing that deep learning can serve as a scalable function approximator consistent with the Kolmogorov–Arnold representation theorem and the universal approximation theorem cited in the literature.

#### ***4.2 Machine Learning and Deep Reinforcement Learning Approaches***

The strong convergence behavior observed in our training procedure is consistent with the deep learning theory discussed in [Goodfellow et al. \(2016\)](#) [11] and the optimization results involving SGD and adaptive variants such as Adam by [Kingma and Ba \(2015\)](#) [42].

Our theoretical statement on SGD convergence corresponds to the class of results also reported in [Bottou, 2018](#) and in reinforcement learning textbooks such as [Sutton and Barto \(2018\)](#) [16, 33].

Compared with existing deep reinforcement learning applications in macroeconomics—such as those by [Fernández-Villaverde et al. \(2023\)](#) [10] and recent applications of neural networks to solve heterogeneous-agent models—our method differs in that:

We focus explicitly on multigenerational OLG structure, a case rarely addressed in DRL macro literature.

We introduce redundant state variables to stabilize the DNN, consistent with insights from representation learning.

We integrate borrowing constraints, adjustment costs, and liquid asset markets, similar in spirit to the economic mechanisms highlighted by [Kaplan et al \(2018\)](#) [35], but embed them inside a machine-learning-driven solution.

Recent works further demonstrate that deep learning and reinforcement learning have reached maturity in handling economic dynamic systems: (1) Deep equilibrium learning in DSGE and heterogeneous-agent models ([Auclert et al., 2020](#)) [36]; (2) Transformer-based approaches for dynamic control ([Zhang et al., 2023](#)) [43]; (3) Stability results for adaptive SGD and Adam in nonconvex environments; (4) RL applications in complex dynamic economies.

These studies provide independent confirmation of the convergence behavior we observe: as the number of episodes increases, stochastic gradient noise diminishes, and the DNN converges toward structural economic decision rules.

Our model's integration of liquid and illiquid assets also matches emerging ML-macro literature emphasizing liquidity channels ([Kaplan et al, 2023](#); [Auclert et al, 2020](#)) [35, 36]. The ability of the DNN to replicate these nonlinear interactions across generations (Figures 4–7) reinforces that modern ML architectures are well suited for high-dimensional incomplete-market environments.

### ***4.3 Evaluation Relative to Analytical Benchmark***

The comparison between analytical and neural network solutions reveals important insights:

(i) Early-generation deviation (Figure 3): The divergence in Generation 1 reflects the instability of the network's early-stage learning and is consistent with optimization noise discussed in [Bertsekas and Tsitsiklis \(1996\)](#) [44] and convergence behavior in stochastic optimization theory. This confirms that without sufficient episodes, the network does not accurately learn the curvature of the value function.

(ii) Mid- and later-generation convergence (Figures 4–7): From Generation 2 onward, the DNN solution converges remarkably well to the analytical benchmark. This mirrors earlier results in approximate dynamic programming ([Powell, 2007](#)) [45], where nonlinear function approximators improve accuracy as long as training data are adequately rich.

The smooth monotonic decision rules obtained by the DNN method also align with theoretical predictions in OLG models with borrowing constraints and capital adjustment costs ([Aiyagari 1994](#); [Kaplan et al. 2018](#); [Huggett 1993](#)) [8, 35, 46].

Recent methodological comparisons indicate that deep learning-based dynamic solvers can match or outperform projection and perturbation methods in high-dimensional settings ([Han and Li, 2023](#)) [47]. These studies highlight several properties that we also observe: (1) Early-stage noise and instability are common when training deep models on stochastic dynamic systems; (2) Monotonicity and concavity of learned policy rules emerge naturally once training converges; (3) High-dimensional value functions are accurately approximated through DNN architectures without explicit discretization.

Thus, the behavior of our model for Generations 2–5 is not only consistent with classical DP theory, but also aligns with modern DL-based dynamic optimization results published between 2021–2025.

### ***4.4 Contribution to the Literature***

Relative to recent post contributions in computational macroeconomics, this study:



(i) Provides one of the first DNN-based scalable frameworks for multigenerational stochastic OLG economies, complementing recent advances in high-dimensional HA/DSGE models.

(ii) Demonstrates that DNNs can recover intergenerational equilibrium decision rules, adding to recent literature on ML-based economy-wide dynamic approximations.

(iii) Extends the ML-macro interface by incorporating multi-asset liquidity frictions, which have recently been highlighted in empirical macro-finance models (Auclert et al., 2020) [36].

This positions the paper alongside the most recent methodological work on neural dynamic equilibrium solvers.

#### **4.5 Limitations and Future Research**

Recent literature suggests several pathways for future research that align with our findings:

(i) Hybrid PINNs—economic models have shown strong convergence for HJB equations, which could stabilize early-generation behavior;

(ii) Attention-based architectures have recently been applied to dynamic control (Zhang et al., 2023) [43] and could enhance representation of life-cycle features;

(iii) RL actor–critic methods with economic constraints are emerging, potentially reducing training cost.

There remains substantial opportunity to extend the DNN-based solution of OLG economies to environments with aggregate uncertainty, disaster-risk shocks, and global incomplete markets—areas of intense research activity post.

### **5. Conclusion**

In this paper, we developed a novel deep neural network (DNN)–based framework for solving dynamic optimization problems in economics, which often involve high-dimensional state spaces, nonlinear objective functions, and complex constraints. The proposed method enables the approximation of optimal policies and value functions without relying on convexity assumptions, thereby overcoming limitations of traditional analytical and numerical techniques.

From a theoretical perspective, we demonstrate the convergence of the learning algorithm based on stochastic gradient descent, even in settings with non-convex loss functions. Empirically, we validate the approach across several standard economic environments—including single-agent consumption–saving models and heterogeneous overlapping-generation structures. Simulation results indicate that DNN-based approximations closely track analytical benchmarks and remain stable across generations and state-space dimensions.

The method shows strong potential for extension to more complex economic models, such as liquidity-constrained economies, heterogeneous-agent macroeconomic models, or large-scale dynamic systems where conventional approaches become computationally infeasible. The use of deep neural networks enhances not only solution accuracy but also computational efficiency and modeling flexibility.

Overall, this study provides a robust, scalable, and versatile framework for solving dynamic optimization problems in economics and opens new avenues for research on convergence, reliability, and large-scale economic modeling.

## REFERENCES

- [1] Bellman, R. (1957). *Dynamic programming*. Princeton University Press.
- [2] Bellman, R. (2010). *Dynamic programming* (2nd ed.). Princeton University Press, 392 pp.
- [3] Lucas, R. E., Jr. (1976). *Econometric policy evaluation: A critique*. Carnegie-Rochester Conference Series on Public Policy, 1(1), 19–46. [https://doi.org/10.1016/S0167-2231\(76\)80003-6](https://doi.org/10.1016/S0167-2231(76)80003-6).
- [4] Sargent, T. J. (1987). *Dynamic macroeconomic theory*. Harvard University Press, 369 pp.
- [5] Judd, K. L. (1998). *Numerical methods in economics*. MIT Press, 656 pp.
- [6] Den Haan, W., and Marcet, A. (1990). Solving the stochastic growth model by parameterized expectations. *Journal of Business and Economic Statistics*, 8(1), 31–34.
- [7] Schmitt-Grohé, S., and Uribe, M. (2004). Solving dynamic equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control*, 28, 755–775. [https://doi.org/10.1016/S0165-1889\(03\)00043-5](https://doi.org/10.1016/S0165-1889(03)00043-5).
- [8] Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109(3), 659–684. <https://doi.org/10.2307/2118417>.
- [9] Krusell, P. and Smith, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5), 868–896. <https://doi.org/10.1086/250034>.
- [10] Fernández-Villaverde, J., Hurtado, S., and Nuño, G. (2023). Financial Frictions and the Wealth Distribution. *Econometric Society: Econometrica*, 91(3), 869-901. <https://doi.org/10.3982/ECTA18180>.
- [11] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press, 800 pp.
- [12] Kelleher, J. D. (2019). *Deep learning*. MIT Press, 296 pp.
- [13] Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5), 359–366. [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8).
- [14] Jiang, Z., et al. (2023). A neural network-based PDE solving algorithm with high precision. *Scientific Reports*, 13, 4479. <https://doi.org/10.1038/s41598-023-31236-0>.
- [15] Sirignano, J., and Spiliopoulos, K. (2018). DGM: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375, 1339–1364. <https://doi.org/10.1016/j.jcp.2018.08.029>.
- [16] Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press, 552 pp.
- [17] Mnih, V., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518, 529–533. <https://doi.org/10.1038/nature14236>.
- [18] Silver, D., et al. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529, 484–489. <https://doi.org/10.1038/nature16961>.
- [19] Maliar, L., Maliar, S., & Winant, P. (2021). Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122, 76–101. <https://doi.org/10.1016/j.jmoneco.2021.07.004>.
- [20] Fernández-Villaverde, J., Guerrón-Quintana, P. (2020). Estimating DSGE models: Recent Advances and Future Challenges. NBER Working Paper No. 27715. <http://www.nber.org/papers/w27715>.
- [21] Fernández-Villaverde, J. (2025). Deep learning for solving economic models (NBER Working Paper No. 34250). <https://doi.org/10.3386/w34250>
- [22] Patrick, C., Jean-Loup, D., & Donatien, H. (2025). Deep learning for continuous-time stochastic control with jumps. *arXiv preprint*. <https://doi.org/10.48550/arXiv.2505.15602>
- [23] Azinovic, M., Luca, J. and Scheidegger, S. (2022). Deep equilibrium nets. *International Economic Review*, 63, 1471-1525. <https://doi.org/10.1111/iere.12575>.
- [24] Amos, B. and Kolter, J. Z. (2021). OptNet: Differentiable optimization as a layer in neural networks. *ICML*. <https://doi.org/10.48550/arXiv.1703.00443>.
- [25] Assouli, M. and Missaoui, B. (2024). Deep policy iteration for high-dimensional mean field games. *arXiv preprint*. [arXiv:2310.10827](https://arxiv.org/abs/2310.10827).

- [26] Kerdabadi, A. E., & Malek, A. (2025). Solving nonlinear and complex optimal control problems via multi-task neural networks. *Scientific Reports*, 15, 25401. <https://doi.org/10.1038/s41598-025-10339-w>.
- [27] Hansen, L. P., and Sargent, T. J. (2001). Robust control and model uncertainty. *American Economic Review*, 91(2), 60–66. <https://doi.org/10.1257/aer.91.2.60>
- [28] Hansen, L. P., and Sargent, T. J. (2003). Robust control of forward-looking models. *Journal of Monetary Economics*, 50(3), 581–604. [https://doi.org/10.1016/S0304-3932\(03\)00026-6](https://doi.org/10.1016/S0304-3932(03)00026-6).
- [29] Bertsekas, D. P. (2005). *Dynamic programming and optimal control*. Athena Scientific, 1, 558 pp.
- [30] Bertsekas, D. P. (2017). *Dynamic programming and optimal control*. Athena Scientific, 2, 576 pp.
- [31] Han, J., Jentzen, A., & E, Weinam. (2018). Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34), 8505–8510. <https://doi.org/10.1073/pnas.1718942115>.
- [32] Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. *Proceedings of COMPSTAT 2010*, 177–186.
- [33] Bottou, L. (2018). *Optimization methods for large-scale machine learning*. Cornell University Press, 95p. <https://doi.org/10.48550/arXiv.1606.04838>.
- [34] Ghadimi, S., and Lan, G. (2013). Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4), 2341–2368. <https://doi.org/10.1137/120880811>
- [35] Kaplan, G., & Violante, G. L. (2018). Microeconomic Heterogeneity and Macroeconomic Shocks. *The Journal of Economic Perspectives*, 32(3), 167–194. <https://www.jstor.org/stable/26473069>.
- [36] Auclert, A., Rognlie, M, and Straub, L, (2020). Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated Hank Model. CEPR Discussion Paper No. DP14279, SSRN. <https://ssrn.com/abstract=3518620>
- [37] Becker, S., et al. (2021). Solving high-dimensional optimal stopping problems using deep learning. *European Journal of Applied Mathematics*, 32, 470–514. <https://doi.org/10.1017/S0956792521000073>.
- [38] Rust, J. (1996). Numerical dynamic programming in economics. *Handbook of Computational Economics*, 1, 619–729. [https://doi.org/10.1016/S1574-0021\(96\)01016-7](https://doi.org/10.1016/S1574-0021(96)01016-7).
- [39] Uribe, M. and Schmitt-Grohé, S. (2017). *Open economy macroeconomics*. Princeton University Press, 648p.
- [40] Chen, M., et al. (2023). Deep reinforcement learning in a monetary model. arXiv preprint. <https://doi.org/10.48550/arXiv.2104.09368>
- [41] Chen, R. (2025). Deep reinforcement learning in a search-matching model of labor market fluctuations. *Economies*, 13(10), 302. <https://doi.org/10.3390/economies13100302>
- [42] Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *ICLR 2015*.
- [43] Zhang, H., et al. (2023). Differentiable optimization layers enhance GNN-based mitosis detection. *Scientific Reports*, 13, 14306. <https://doi.org/10.1038/s41598-023-41562-y>.
- [44] Bertsekas, D. P., & Tsitsiklis, J. (1996). *Neuro-dynamic programming*. Athena Scientific, 506pp.
- [45] Powell, W. B. (2007). *Approximate dynamic programming: Solving the Curses of Dimensionality*. Wiley. <https://doi.org/10.1002/9780470182963>.
- [46] Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5–6), 953–969. [https://doi.org/10.1016/0165-1889\(93\)90024-M](https://doi.org/10.1016/0165-1889(93)90024-M).
- [47] Han, Y., & Li, N. (2023). A deep neural network algorithm for multiple stopping. *Communications in Nonlinear Science and Numerical Simulation*, 117, 106881. <https://doi.org/10.1016/j.cnsns.2022.106881>.