

Peiyuan Zhang

Academic Job Market 2026–2027

Ph.D. Candidate in Economics, School of Economics, Peking University
Beijing, China | zpyuan@pku.edu.cn | cdjnzpy.github.io

EDUCATION

Peking University, School of Economics 2022–Present
Ph.D. in Economics; expected June 2027.

Peking University, School of Economics 2018–2022
Bachelor of Economics.

RESEARCH AND TEACHING FIELDS

Research. Macroeconomics; Computational Economics; AI for Economics; Public Economics. Focus on AI solvers for dynamic economic models, heterogeneous-agent economies, and multi-agent economic simulation.

Teaching preparation. Macroeconomics and computational economics; machine learning, deep learning, and reinforcement learning; courses connecting AI methods and economic analysis.

RESEARCH

Working Papers

Informed Actor-Critic Method for Optimal Control in Economic Dynamics.

with Kenneth L. Judd, Bo Li, and Serguei Maliar. [Paper]

Job Market Paper

This paper applies deep reinforcement learning to high-dimensional optimal control in dynamic economic models. We introduce an informed actor–critic (iAC) framework that embeds structural equations and constraints directly into the policy optimization loop. Unlike standard model-free reinforcement learning algorithms that treat the economy as a black box, the iAC method leverages analytical economic transitions. Computational applications to benchmark consumption–savings, Krusell–Smith, and overlapping-generations economies show that iAC reduces the need for exploration noise and provides a precise, scalable, and sample-efficient framework for complex macroeconomic control problems.

Solving Discrete-Continuous Dynamic Choice Models: A Soft-to-Hard AI Solver with an Application to Sovereign Default.

with Bo Li and Serguei Maliar. [Paper]

To solve discrete-continuous dynamic choice models, the quantitative literature has long relied on effective permanent structural smoothing to handle non-differentiable kinks. While traditional methods produce robust qualitative insights, their permanent modifications can alter quantitative predictions. To solve the true, unsmoothed economy, we introduce a soft-to-hard (StH) deep reinforcement learning framework. By systematically annealing temporary smoothing to zero, StH enables gradient propagation across decision boundaries before enforcing hard-switch operators. We implement this framework in two complementary ways. The grid-based StH-Q implementation delivers high numerical precision in the low-dimensional benchmark, whereas the continuous-state, continuous-action StH-AC implementation avoids storing a full state–action table and provides a scalable route to higher-dimensional applications.

W-exp: A Routed Post-Decision Expectation Adapter for Dynamic Economic Solvers.

with Bo Li and Serguei Maliar. [Paper]

This paper proposes W-exp, an amortized expectation module that converts simulated shock information into a reusable learned conditional expectation function over post-decision states. W-exp is algorithm-agnostic and improves the expectation-construction step shared by neural and simulator-based solvers. Under matched transition-sample budgets, W-exp lowers final policy-function error relative to Monte Carlo benchmarks by 76.2% with Q-learning, 31.4% with informed actor-critic, and 26.9% with soft actor-critic in a consumption–saving model. In a two-shock Krusell–Smith economy, a solver using the learned exact-label variant attains 42.9% lower policy-function error than the same solver using direct exact integration.

Gradient Descent of the Discrete-Time Ramsey Growth Model.

with Bo Li, Yapeng Qian, and Shenghao Zhu.

This paper provides an optimization-oriented study of the discrete-time Ramsey growth model, developed as a companion to our continuous-time analysis, replacing every continuous-time device by its discrete-time counterpart. We derive the discrete Euler–Lagrange residual as a functional gradient and use gradient-based procedures

to validate the Ramsey dynamics: strict concavity and global optimality from the second variation, local strong convexity from a discrete Poincaré inequality, and exponential convergence of the sequence-space gradient flow from a Lyapunov–Grönwall argument. We construct a quadratic Lyapunov function and show, from strict monotonicity of the optimal capital policy rather than by linearization, that the optimal policy never overshoots, certifying global convergence to the steady state. We then relate the resulting schemes to deep deterministic policy gradient methods, identifying exact analytical counterparts for residual relaxation, target networks, the actor–critic split, and off-policy sampling.

Gradient Descent of Ramsey Growth Model.

with Bo Li, Yapeng Qian, and Shenghao Zhu.

This paper provides an optimization-oriented study of the continuous-time Ramsey growth model, connecting classical stability analysis with residual-minimization approaches for solving Hamilton–Jacobi–Bellman (HJB) equations. We revisit the planner’s primal formulation through a variational perspective and validate the Ramsey dynamics numerically. We construct a Lyapunov function and show that it decreases along saddle-point stable trajectories, certifying convergence to the steady state. We then formulate a residual loss functional induced by the HJB operator and study its analytical properties, including well-posedness, regularity, and convexity in an infinite-dimensional setting, enabling a principled comparison of first-order and second-order optimization schemes. We extend the construction to deterministic and stochastic environments, discussing how diffusion terms modify the HJB residual.

Contribution Base Constraints and Flexible Retirement: Policy Trade-offs and Combination Effects in China’s Pension Reform.

with Bo Li and Yi Lu.

We build a heterogeneous-agent overlapping-generations model with an endogenous retirement choice to quantify the separate and joint effects of contribution base adjustments and flexible retirement, specifying contribution base bounds, contribution histories, and the benefit rules of the two pension accounts in a single household problem. The lower contribution bound is the main policy margin: lowering it relieves low-income contributors but reduces replacement rates and widens the pension deficit, with limited aggregate gains. Flexible retirement raises output and improves the pension balance through a small number of marginal individuals with high current productivity and low historical contribution bases. Combining the two policies preserves the relief effect while offsetting much of the benefit and financing pressure. A Shapley decomposition attributes relief mainly to the lower bound, and the output and financing gains mainly to flexible retirement. Shifting the contribution window forward to early in the working life reduces the deficit improvement from 4.65% to 0.58%.

AI Taxing AI: A Multi-Agent Reinforcement Learning Approach to Optimal Taxation.

with Bo Li. [Paper]

Artificial intelligence (AI) is improving productivity, yet its economy-wide diffusion may reshape the labor market and the distribution of income. To address these issues, we develop a Multi-Agent Reinforcement Learning (MARL) model featuring endogenous skill accumulation. We find that while AI boosts aggregate output through rapid skill sharing, it displaces middle-skilled human workers and increases inequality. However, optimal taxation can redistribute the technological surplus to achieve Pareto improvements. To find the optimal tax schedule, we employ Dual-Loop Multi-Agent Reinforcement Learning to account for both the government’s fiscal policy and agents’ behavior. Furthermore, we demonstrate that a “Digital Commons” regime—combining universal access with real-time knowledge sharing—maximizes social welfare by mitigating the trade-off between efficiency and equity.

Forecasting Crude Oil Spot Prices Using a Transformer-BiLSTM Architecture with NRBO-Based Hyperparameter Optimization.

with Wenhui Huang et al. | Revise and resubmit, *Sustainable Futures*.

Combines Transformer and BiLSTM architectures with hyperparameter optimization for Brent and WTI spot-price forecasting.

Work in Progress

Reduced-State Deep Learning for Heterogeneous Agent Models.

A reduced-state neural solution method that approximates the distributional state using a low-dimensional vector of moments, keeping heterogeneous-agent models with aggregate uncertainty computationally tractable.

Emergence of Universal Equivalents with Deep Reinforcement Learning.

A multi-agent reinforcement learning environment in which agents evolve from autarky to barter and converge on a universal equivalent, allowing price discovery and monetary emergence to be studied without strong institutional assumptions.

Publications

Asset Pricing in China's Stock Market.

with Shuaiyu Jiang | *Finance & Economics Vision*, 1(1), 2022. [Paper]

This paper uses traditional machine learning methods and deep neural networks based on both firm-specific characteristics and macroeconomic variables to price China's A-share stock market. We give the stochastic discount factor a flexible form and compare different models' performances. Since the Chinese government adopts various policies to maintain financial stability, we borrow the idea from generative adversarial networks to find the true SDF by selecting moment conditions that minimize return volatility.

TEACHING EXPERIENCE

Principles of Macroeconomics (undergraduate)

Spring semesters, 2022–2026

Teaching Assistant, Peking University

- Assisted with course coordination and recitation teaching; contributed to the design of problem sets, computational exercises, and midterm and final examination materials.
- Developed an interactive lesson script and accompanying Jupyter lab on the Solow model, integrating model derivations, parameter experiments, visualization, and policy discussion.

PKU–Zurich PhD Summer School on Machine Learning for Macroeconomics and Finance

July 2026

Teaching Support for Serguei Maliar's hands-on sessions

- Prepared and maintained Chinese–English guides for environment setup and notebook execution in TensorFlow, PyTorch, and JAX.
- Assisted with guided exercises, notebook execution, and troubleshooting, including Krusell–Smith materials.

Computational Economics (graduate)

Spring 2022

Teaching Assistant, Peking University

Artificial Intelligence and Society (undergraduate)

Spring 2022

Teaching Assistant, Peking University

- Supported teaching on AI fundamentals and its social implications, including interdisciplinary discussions.

COURSES PREPARED TO TEACH

Economics core. Principles of Macroeconomics; Intermediate and Advanced Macroeconomics; Computational Economics.

AI foundations. Machine Learning for Economists; Deep Learning; Reinforcement Learning.

Economics × AI. AI for Solving Dynamic Economic Models; Multi-Agent Learning and Economic Simulation; AI, Distribution, and Public Policy.

RESEARCH PROJECT PARTICIPATION

Role in each project: Core participant.

NSFC, Original Exploration Program. Construction of Macroeconomic Models Based on Intelligent Algorithms and Agents and Their Application in Policy Evaluation.

NSFC, Young Scientists Fund Program. Design of Optimal Property Tax Reform: A Quantitative Study Based on Heterogeneous Household Models.

Peking University / National Intelligent Social Governance Experimental Base (Donghu High-Tech Zone), Open Project. Research on Macroeconomic Theories, Algorithms, and Models Based on AI—Scenario Simulation in Donghu High-Tech Zone.

New Engineering Interdisciplinary Youth Project. Emergence and Evolution of Currency Based on Multi-Agent Learning.

ACADEMIC SERVICE

Member of the organizing team for the 2026 SMLE Annual Conference and the PKU–Zurich PhD Summer School on Machine Learning for Macroeconomics and Finance: participant communication, schedule coordination, teaching-code repository maintenance, and on-site support.

SKILLS

Computing. Python (PyTorch), Fortran, C/C++, CUDA, MATLAB; OpenMP.

Methods. Dynamic programming, numerical optimization, automatic differentiation, deep learning, reinforcement learning, multi-agent simulation, parallel and GPU computing.

Modeling experience. Huggett, Aiyagari, Krusell–Smith, HANK, OLG, and sovereign default models.

Languages. Chinese (native); English (academic communication and teaching).

REFERENCES

Prof. Bo Li School of Economics, Peking University
boli.econ@pku.edu.cn

Prof. Serguei Maliar Department of Economics, Santa Clara University
maliars@stanford.edu