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. Computational Macroeconomics; Artificial Intelligence; Public Finance.

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 | **Prepared for *Econometrica*.** [Paper]

Job Market Paper

Abstract: 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.

Discrete-Continuous Dynamic Choice: A Soft-to-Hard Reinforcement Learning Framework with an Application to Sovereign Default.
with Bo Li and Serguei Maliar | **Under review at *Journal of Monetary Economics*.** [Paper]

Abstract: Discrete-continuous dynamic choice models create non-differentiable kinks for gradient solvers, while permanent structural smoothing can alter quantitative predictions. We introduce a soft-to-hard (StH) deep reinforcement learning framework that anneals temporary smoothing to zero, enabling gradient propagation across decision boundaries before enforcing the model's true hard-switch operators. The grid-based StH-Q implementation delivers high numerical precision in a low-dimensional sovereign-default benchmark, whereas continuous-state, continuous-action StH-AC 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]

Abstract: This paper proposes W-exp, an algorithm-agnostic amortized expectation module that converts simulated shock information into a reusable conditional-expectation function over post-decision states. Under matched transition-sample budgets in a consumption–saving model, 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 two-shock Krusell–Smith economy, the learned exact-label variant attains 42.9% lower policy-function error than the same solver using direct exact integration.

DSPG: Learning Distribution-Conditioned Policies for Heterogeneous-Agent Economies.
with Chiyuan Wang and Bo Li | **Submitted to ICLR 2027.**

Abstract: Household heterogeneity makes the cross-sectional distribution an infinite-dimensional recursive state: choices move the distribution, which feeds back through equilibrium prices. We introduce distribution-based structural policy gradients (DSPG), conditioning household policies directly on the asset–income histogram while the model supplies budget accounting, market clearing, and distribution transitions. A one-dimensional U-Net exploits the ordered asset grid, retains the full histogram without a prespecified moment map, and uses only 783 parameters. In a Huggett economy with aggregate shocks, DSPG produces smooth equilibrium policies, small asset-market residuals, and the lowest Euler diagnostic among four modern benchmark algorithms.

Gradient Descent of the Discrete-Time Ramsey Growth Model.
with Bo Li, Yapeng Qian, and Shenghao Zhu.

Abstract: Provides an optimization-oriented analysis of the discrete-time Ramsey growth model. We derive the Euler–Lagrange residual as a functional gradient and establish global optimality, local strong convexity, and exponential convergence using second-variation, discrete Poincaré, and Lyapunov arguments. A quadratic Lyapunov function and strict monotonicity of the optimal capital policy establish non-overshooting and global convergence without linearization. The resulting schemes reveal analytical counterparts of residual relaxation, target networks, the actor–critic split, and off-policy sampling in deep deterministic policy-gradient methods.

Gradient Descent of Ramsey Growth Model.
with Bo Li, Yapeng Qian, and Shenghao Zhu.

Abstract: Connects variational stability analysis of the continuous-time Ramsey model with residual minimization for Hamilton–Jacobi–Bellman equations. A Lyapunov function certifies convergence along saddle-path stable trajectories. We then formulate the HJB-induced residual loss and study its well-posedness, regularity, and convexity in an infinite-dimensional setting, enabling a principled comparison of first- and second-order optimization methods. The construction extends to deterministic and stochastic environments, clarifying 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 | **Prepared for *Economic Research Journal* (*Jingji Yanjiu*).**

Abstract: Builds a heterogeneous-agent OLG model with endogenous retirement to quantify the separate and joint effects of contribution-base adjustments and flexible retirement. The model jointly represents contribution-base bounds, contribution histories, and benefit rules for two pension accounts. Lowering the contribution floor relieves low-income contributors but reduces replacement rates and widens the pension deficit. Flexible retirement raises output and improves pension finances through marginal workers with high current productivity and low historical contribution bases. A Shapley decomposition clarifies how the reforms divide the roles of burden relief, output expansion, and fiscal improvement; moving the contribution window toward early working life sharply weakens the deficit reduction.

RESEARCH (CONTINUED)

Working Papers (continued)

AI Taxing AI: A Multi-Agent Reinforcement Learning Approach to Optimal Taxation.
with Bo Li. [Paper]
Abstract: Develops a multi-agent reinforcement learning economy with endogenous skill accumulation to study AI diffusion, labor displacement, inequality, and optimal taxation. AI raises aggregate output through rapid skill sharing but displaces middle-skilled workers. A dual-loop algorithm jointly accounts for fiscal policy and behavioral responses, while a *Digital Commons* regime combines universal access with real-time knowledge sharing to improve the efficiency–equity trade-off.

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*.** [Paper]
Abstract: Develops a hybrid forecasting pipeline for Brent and WTI crude-oil spot prices. Transformer attention captures long-range dependencies, BiLSTM layers model bidirectional temporal patterns, and NRBO-based search tunes the architecture and training hyperparameters across multiple forecast horizons.

Work in Progress

Reduced-State Deep Learning for Heterogeneous Agent Models.
Abstract: Develops a reduced-state neural solution method that replaces the full cross-sectional distribution with a low-dimensional vector of moments. The representation enters household-policy and aggregate-law-of-motion networks, reducing the state dimension while retaining the distributional feedback needed to solve heterogeneous-agent economies with aggregate uncertainty.

Emergence of Universal Equivalents with Deep Reinforcement Learning.
Abstract: Constructs a multi-agent reinforcement learning environment in which agents evolve from autarky to barter and converge on a universal equivalent. The framework studies endogenous price discovery, specialization, and monetary emergence under decentralized learning, without imposing a medium of exchange or strong institutional assumptions in advance.

Publication

Asset Pricing in China's Stock Market.
with Shuaiyu Jiang | *Finance & Economics Vision*, 1(1), 2022. [Paper]
Abstract: Uses firm-specific characteristics and macroeconomic variables with traditional machine learning methods and deep neural networks to price China's A-share market. We give the stochastic discount factor a flexible form and compare models across predictive and pricing performance. Motivated by the policy-rich environment of the Chinese market, an adversarial procedure selects moment conditions that minimize return volatility and helps identify the stochastic discount factor.

TEACHING AND CURRICULUM DEVELOPMENT

Principles of Macroeconomics (undergraduate), Peking University Spring 2022–2026
Teaching Assistant

- Recitation teaching and course coordination; designed problem sets, computational exercises, and examination materials.
- Developed an interactive Solow-model lesson and Jupyter lab linking 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 bilingual setup and notebook guides and supported exercises and troubleshooting in TensorFlow, PyTorch, and JAX.

Computational Economics (graduate), Peking University Spring 2022
Teaching Assistant

Artificial Intelligence and Society (undergraduate), Peking University Spring 2022
Teaching Assistant; supported teaching on AI fundamentals, social implications, and interdisciplinary discussion.

Digital Textbook Development: *Principles of Macroeconomics* 2026–Present
Project participant; contributing to content authoring and the integration of chapter slides, interactive code, data visualization, and assessment resources.

Open Teaching and Learning Resources Public reading notes, slides, and replication code on computational macroeconomics, heterogeneous-agent models, and neural solution methods.

COURSES PREPARED TO TEACH

Economics. Principles, Intermediate, and Advanced Macroeconomics; Computational Economics; Quantitative Public Economics and Social Security.

AI for economists. Machine Learning; Deep Learning; Reinforcement Learning.

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

RESEARCH PROJECT PARTICIPATION

Core participant in each project.

NSFC Original Exploration Program. Macroeconomic models based on intelligent algorithms and agents for policy evaluation.

NSFC Young Scientists Fund Program. Optimal property-tax reform using heterogeneous-household models.

Peking University / National Intelligent Social Governance Experimental Base. AI-based macroeconomic theories, algorithms, and scenario models for Donghu High-Tech Zone.

New Engineering Interdisciplinary Youth Project. Emergence and evolution of currency through multi-agent learning.

ACADEMIC SERVICE

Organizing team member for the 2026 SMLE Annual Conference and the PKU–Zurich PhD Summer School: participant communication, schedule coordination, teaching-code repository maintenance, and on-site support.

SKILLS

Computing. Python (PyTorch), Fortran, C/C++, CUDA, MATLAB, OpenMP; parallel and GPU computing.

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