Growth and (Re)distribution in China and Korea

Bayesian Estimation of a Non-Stationary Heterogeneous-Agent Model

Wending Liu

Australian National University

AI for Economics and Finance Conference, University of Turin

August 2026

Motivation and Overview

- Motivation
  - Growth miracles in China and Korea
  - Growth-accounting and developmental-state explanations
  - The interaction between growth and (re)distribution
- Modelling strategy
  - A non-stationary heterogeneous-agent model
  - A Bayes-net representation of the model
  - Agents with bounded rationality
- Computational method
  - Amortized optimization and inference with deep learning
  - Supervised and unsupervised learning
- Counterfactual experiments
  - How TFP catch-up, education expansion, and demography shape growth and distribution
  - The roles of state capital and fiscal efficiency
  - Potential sources of inequality

Model at a Glance

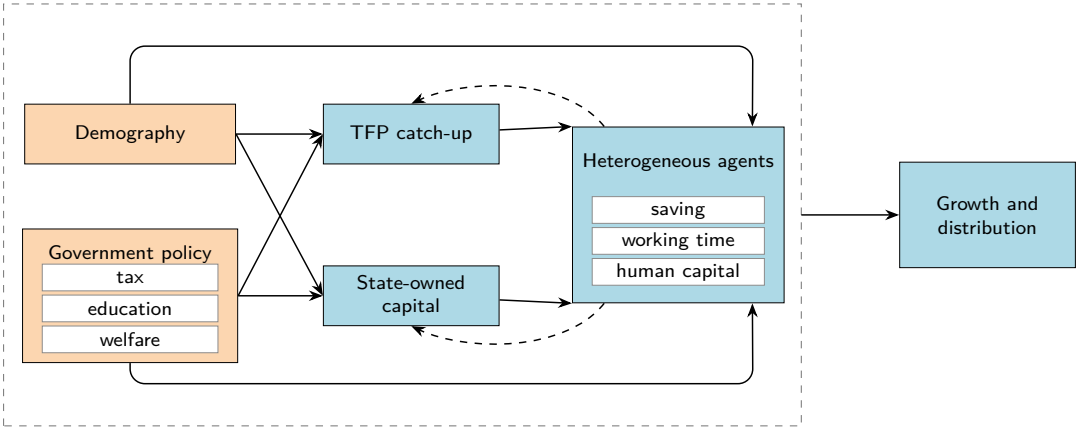

Model Features

- Non-stationary environment
  - Time-varying TFP, demography, and government policy
  - Time-dependent decision rules
- Rich heterogeneity and uncertainty
  - Idiosyncratic labor-productivity risk + TFP risk
  - Family-specific random preference parameters
- Nested CES production
  - Private capital + state capital $\rightarrow$ effective capital
  - Effective capital + effective labor $\rightarrow$ output
- Human-capital-driven technology catch-up
  - Learning by doing $\rightarrow$ endogenous human capital (Imai and Keane, 2004)
  - Education-weighted human-capital density $\rightarrow$ faster catch-up
- Multiple roles of government
  - Tax collector
  - State-capital owner and investor
  - Provider of education and the social safety net

Certainty-Equivalent (CE)-Anchored Quasi-Rational Equilibrium

- Temporary-equilibrium concept (Grandmont, 1977; Adam and Marcet, 2011)
  - Agents optimize under subjective beliefs about future prices
  - Avoids Bellman equations with infinite-dimensional states under RE (Moll, 2026)
- Finite-window CE forecast
  - Future primitive innovations are fixed at their CE values
  - The price anchor clears markets in the shock-free imagined economy
- Subjective price risk
  - The subjective price process is centered on the price anchor

Agents' Decision Rules: One-Step Approximate Rollout

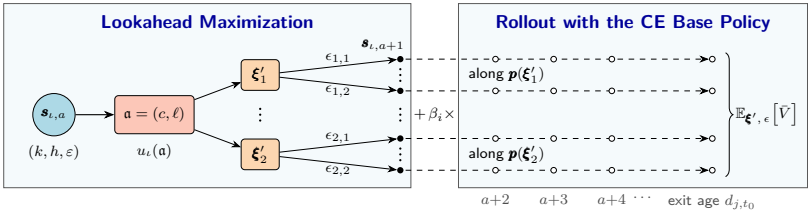

$$\hat{\pi}_{\iota,a}(\boldsymbol{s}) \in \arg\max_{\mathfrak{a}\in\mathfrak{A}_{\iota,a}} \left\{ u_\iota(\mathfrak{a}) + \beta_i \iint \bar{V}_{\iota,a+1}(\boldsymbol{s}';\boldsymbol{p}(\boldsymbol{\xi}'))\hat{p}_{\iota,t}(\boldsymbol{s},\mathfrak{a},d\boldsymbol{s}' \mid \boldsymbol{\xi}')\, dF_{\boldsymbol{\xi}}(\boldsymbol{\xi}') \right\}.$$

- Agent ι, at age a and in state $\boldsymbol{s}$, chooses optimal consumption and labor supply: $\hat{\pi}_{\iota,a}(\boldsymbol{s})$
  - $\bar{V}_{\iota,a+1}$: Continuation value under the CE base policy
- Outer integral
  - $\boldsymbol{p}(\boldsymbol{\xi}')$: Future price path
  - Price-risk innovation $\boldsymbol{\xi}'$ tilts the CE price path
- Inner integral
  - Next-period state $\boldsymbol{s}' \sim \hat{p}_{\iota,t}(\cdot \mid \boldsymbol{\xi}')$, driven by the labor-productivity innovation ε'

Computation Pipeline

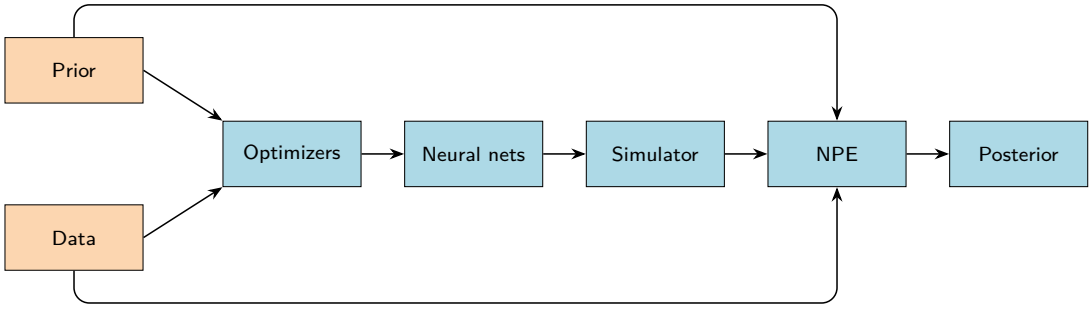

Estimation needs: $50{,}000$ draws $\times$ 2 countries $\times$ 100 MC families $\times$ 21–49 cohorts/family $\times$ 38/58 years (CN/KR)

< 20 **h**
generate labels and simulations

< 3 **h**
train all networks

seconds
posterior at observed data

Amortizing Agents' Decision Rules

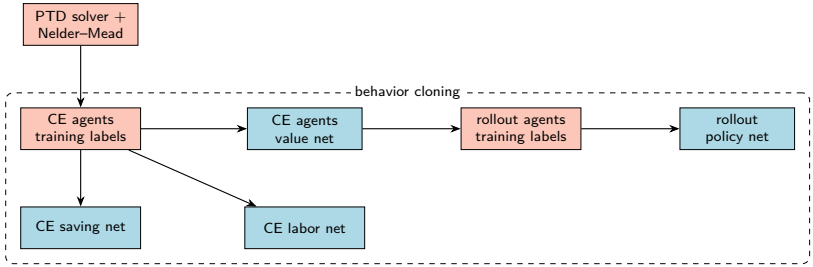

- Expert labels
  - PTD algorithm (Liu and Iskhakov, 2025) solves the saving subproblem;
  - Nelder–Mead searches the working-time path
  - 1 million CE labels + 100,000 exact-rollout labels
- Held-out accuracy
  - Saving-rate MAE: 0.013; working-time MAE: 0.011 per year
  - Value-function R^2 : 0.999
  - Median rollout value gap: 0.05%; 94.5% below 1%

The Policy-Tree Descendant (PTD) Algorithm

$d-1$

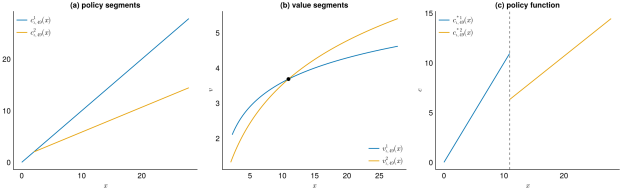

$d-2$

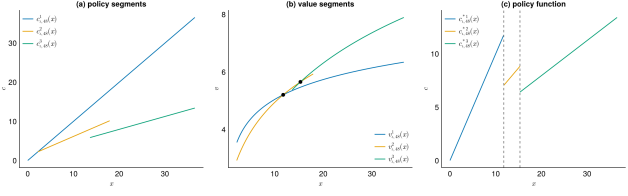

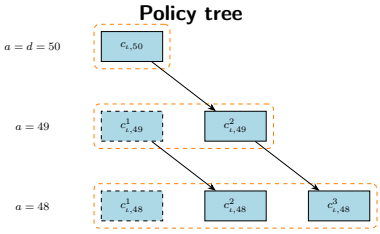

- The policy tree recursively traces affine policy segments.
- Intersections identify the kinks at which the optimal policy switches segments.

Objective-Based Training of the CE Price-Anchor Network

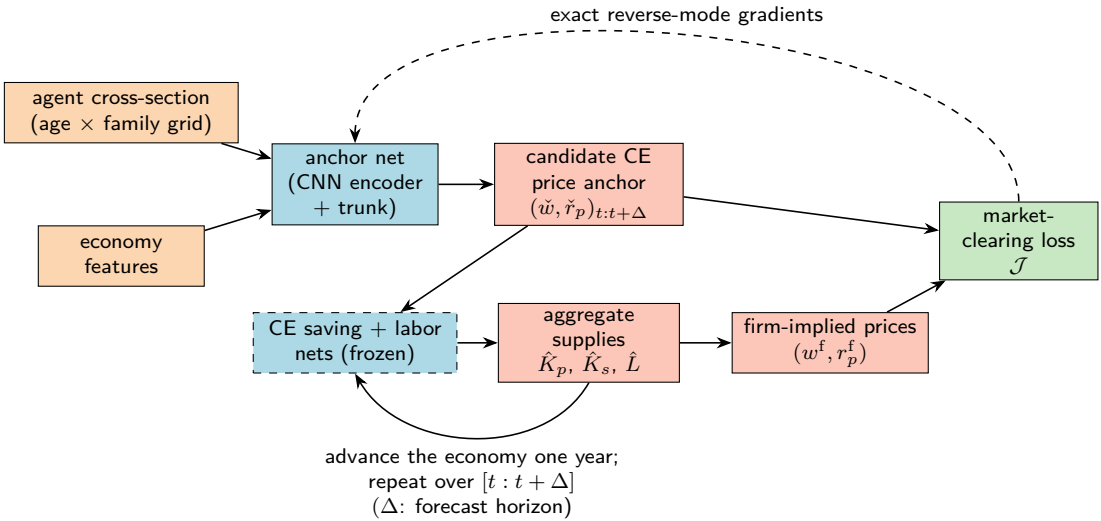

From Price-Anchor Training to Simulator Deployment

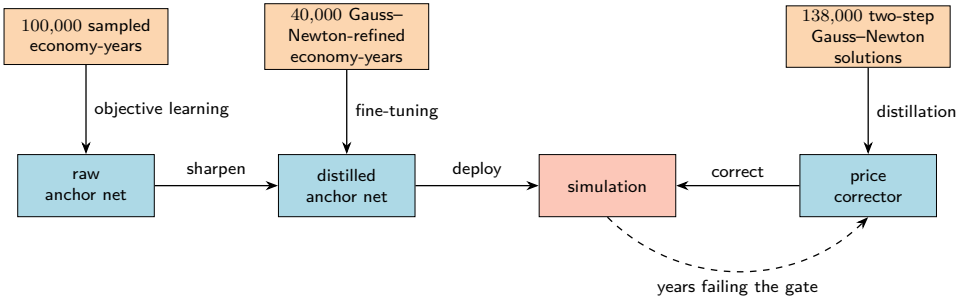

- The price-anchor network is distilled before deployment and corrected during simulation.
- Current states $\rightarrow$ rollout policy $\rightarrow$ consumption–labor choices $\rightarrow$ budget update $\rightarrow$ realized prices $\rightarrow$ updated agent states $\rightarrow \cdots$

Neural Posterior Estimator (NPE)

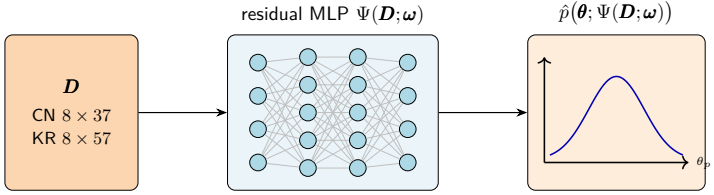

$$\boldsymbol{\omega}^\star = \arg\min_{\boldsymbol{\omega}} \int \left[\int p(\boldsymbol{\theta} \mid \boldsymbol{D}) \log \frac{p(\boldsymbol{\theta} \mid \boldsymbol{D})}{\hat{p}(\boldsymbol{\theta}; \Psi(\boldsymbol{D};\boldsymbol{\omega}))} \, d\boldsymbol{\theta} \right] p(\boldsymbol{D}) \, d\boldsymbol{D}.$$

- Neural posterior mapping
  - $\boldsymbol{D}$: data; $\boldsymbol{\theta}$: model parameters; $\boldsymbol{\omega}$: network weights
  - Neural Network $\Psi(\boldsymbol{D};\boldsymbol{\omega})$ parameterizes posterior $\hat{p}(\boldsymbol{\theta} \mid \boldsymbol{D})$
- Expected KL objective
  - Inner integral: posterior gap for one dataset
  - Outer integral: average gap over model-generated datasets
- Monte Carlo training
  - Prior draw $\boldsymbol{\theta}_i \rightarrow$ simulated data $\boldsymbol{D}_i \rightarrow$ minimize neural posterior NLL

Posterior Predictive Check

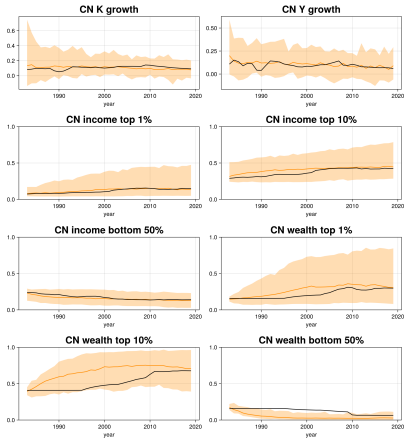

China, 1982–2019

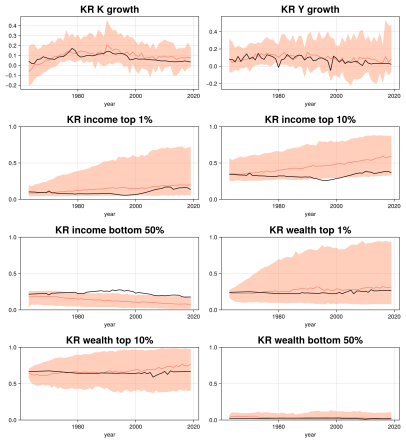

Korea, 1962–2019

Observed data are shown in black; the posterior-predictive median and 5–95% predictive band are shown in color.

The Effect of TFP Catch-Up

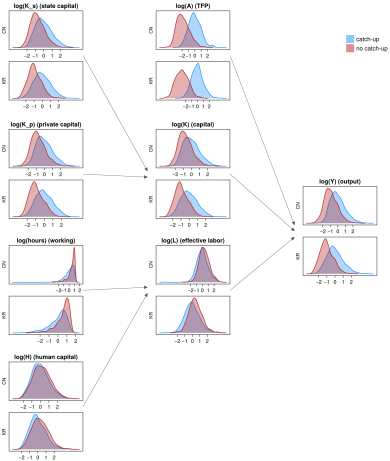

Axes are in each panel's baseline standard deviations; 0 marks the baseline mean.

- 2019 posterior median (China/Korea)
  - TFP: $\uparrow 135\%/195\%$
  - Output per capita: $\uparrow 217\%/306\%$
- Growth through productivity and capital
  - Productivity $\uparrow$ + capital deepening
  - Working time $\downarrow$ + human capital $\downarrow$
- Wealth-distribution effect: ambiguous
  - Wage and capital return channels offset

The Effect of the Education Expansion

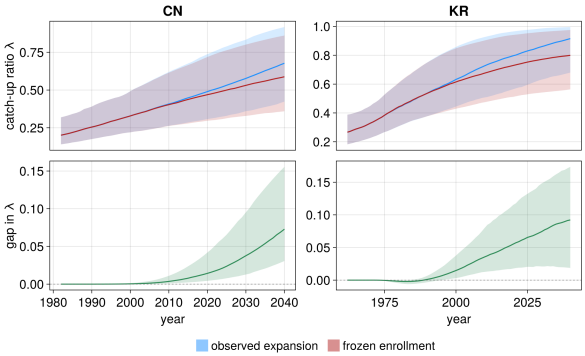

λ: TFP relative to the country-specific ceiling.

- Freeze high-school + college entry shares at 1982/1962 levels (CN/KR)
- Output contribution (median; CN/KR)
  - 2019: $\uparrow 3.6\%/16.6\%$
  - 2040: $\uparrow 25.7\%/39.9\%$
  - Only 2040 bands exclude zero
- Why is the 2040 effect larger?
  - Expansion cohorts gain workforce share
  - China's mass college expansion: 1999 onward
  - Far from frontier: modest human capital density suffices
  - Near frontier: greater human capital density is needed

The Impact of State-Owned Capital

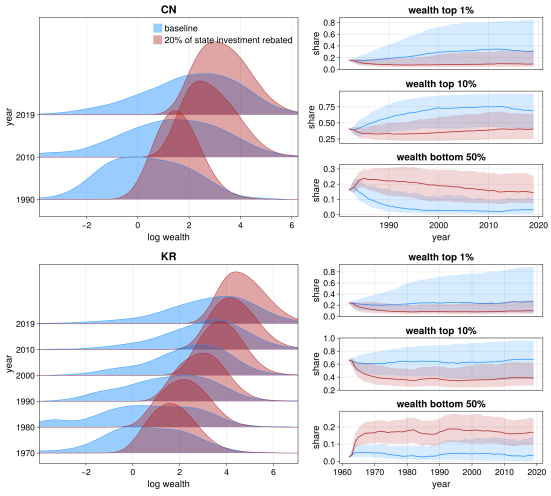

- Redirect 20% of annual state investment uniformly to private-capital
- 2019 median (CN/KR)
  - State capital: $\downarrow 49.1\%/63.8\%$
  - Output: $\downarrow 22.0\%/26.2\%$
  - Working time $\downarrow$
  - Learning by doing $\downarrow$; human capital $\downarrow$
- Suggested growth–inequality trade-off
  - Repeated credits compound household wealth
  - Top-10%: $\downarrow 23.4/24.3$ pp
  - Bottom-50%: $\uparrow 10.0/11.8$ pp

The Impact of Fiscal Efficiency

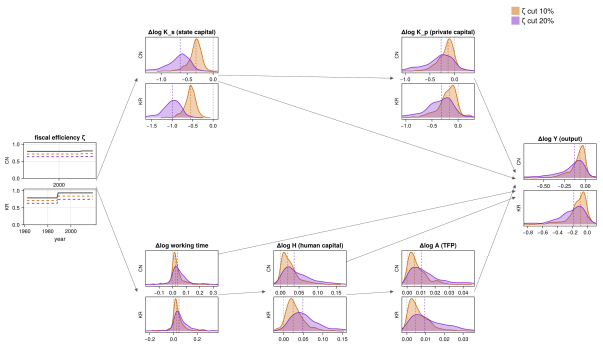

2019 log differences; gray: zero; dashed: medians.

- Fiscal efficiency $\downarrow 10\%/20\%$
  - Lost resources leave the economy
  - No household transfer
- 10% cut in 2019 (median; CN/KR)
  - State capital: $\downarrow 27.9\%/42.2\%$
  - Output: $\downarrow 5.6\%/8.9\%$
- Median adjustment across both cuts
  - State/private capital $\downarrow$ (complements)
  - Working time: $\uparrow 1.9\%$–4.7%
  - Human capital/adoption speed: $\uparrow$
  - TFP: $\uparrow 0.5\%$–1.0%

The Impact of a More Generous Social Safety Net

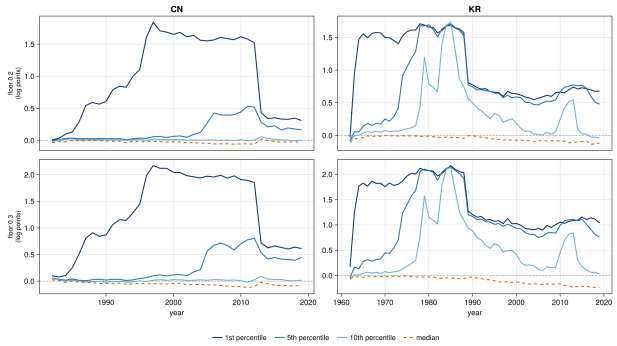

Median log differences; upper/lower: floor $0.2/0.3$.

- Counterfactual
  - Floor/wage ratio: 0.2 or 0.3
- 2019 output (median; CN/KR)
  - Floor 0.2: $\downarrow 3.2\%/7.6\%$
  - Floor 0.3: $\downarrow 10.0\%/19.3\%$
- Consumption incidence at a 0.2 floor
  - First percentile: $\uparrow 36\%/97\%$
  - Tenth percentile: ≈ 0; median: $\downarrow$
- Median growth–equity trade-off
  - Bottom consumption $\uparrow$; state capital $\downarrow$
  - At 0.3, private capital also $\downarrow$

Other Potential Sources of Growth and Inequality

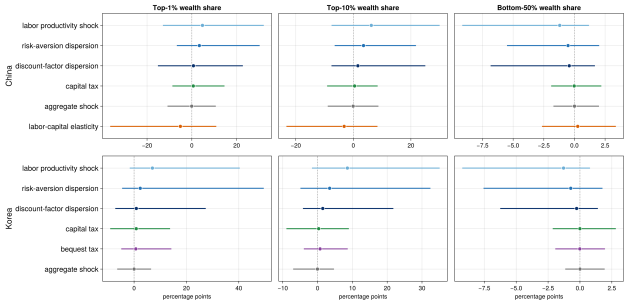

One-source-off wealth-share effects: 2019 medians (pp); all 90% intervals include zero.

- Fertility decline: 2019 median (CN/KR)
  - $\text{Pop} \downarrow \rightarrow Y/\text{Pop} \uparrow 5.7\%/9.8\%$
  - Aggregate output/wealth shares: no clear effect
- Top-10% marginal contributions (median pp)
  - Labor shock: $6.2/8.6$ (largest tested)
  - Risk aversion: $3.5/3.5$
  - Discount factor: $1.6/1.5$
- Other marginal contributions (median)
  - China: labor–capital complementarity < 0
  - TFP shock $+$ tested taxes: near zero

Conclusion

- **Computational Framework**
  - Non-stationary heterogeneous-agent model of growth and distribution
  - CE-anchored quasi-rational equilibrium + amortized optimization with deep learning
  - Bayesian estimation using China and Korea data
  - Counterfactual experiments as interventions on the Bayes net
- **Main findings**
  - TFP catch-up: dominant growth driver; distribution effect unclear
  - Education/fertility decline: per capita output ↑; small distribution effects
  - State investment / safety nets: growth–equality trade-offs
  - Labor-productivity shock: larger marginal effect on inequality than preference dispersion
- **Limitations**
  - Closed economy + one asset
  - Omits housing/land, and other country specific institutions
- **Future research**
  - Open economy + second asset + richer institutional channels