

Growth and (Re)distribution in China and Korea: Bayesian Estimation of a Non-Stationary Heterogeneous Agent Model*

Wending Liu[†]

August, 2026

Abstract

I analyze growth and (re)distribution jointly in China (1982–2019) and Korea (1962–2019) inside a non-stationary heterogeneous-agent model. Two innovations make Bayesian estimation feasible: a certainty-equivalent-anchored quasi-rational equilibrium concept which replaces full rational expectations with a finite-horizon deterministic forecast; and a deep-learning pipeline that amortizes the agent’s decision problem, the equilibrium computation, and the posterior itself. Counterfactual experiments find that TFP catch-up is the dominant driver of output growth yet leaves the wealth distribution untouched, that the education expansion and the fertility decline likewise buy per-capita growth with little distributional footprint, that state investment promotes growth while potentially concentrating wealth, that a thin social safety net buys growth at a large cost to consumption at the bottom, and that among the model’s estimated primitives idiosyncratic labor-productivity risk has the largest marginal impact on wealth inequality.

Keywords: growth and distribution, non-stationary heterogeneous agent model, certainty-equivalent-anchored quasi-rational equilibrium, Bayesian estimation, likelihood-free inference, amortized optimization, deep learning

*I am indebted to my supervisor Fedor Iskhakov for his continuous guidance and support on this paper. I acknowledge helpful comments from Tim Kam, Tue Gørgens, Juergen Meinecke, and John Stachurski. I especially thank Minhee Chae for her good suggestions and for her help with the Korean data.

[†]Australian National University Wending.Liu@anu.edu.au

1 Introduction

After the Second World War, many late-industrializing economies in Asia, Latin America, and Africa attempted to catch up with the already-industrialized West, but only a handful of them, concentrated in East and Southeast Asia, succeeded. China and Korea are two of the most striking examples.¹ In China, the Twelfth National Congress of the Chinese Communist Party proposed socialism with Chinese characteristics in 1982; from then to 2019, per-capita real GDP grew from \$868 to \$18,358. In Korea, the First Five-Year Economic Development Plan launched the state-directed industrialization drive in 1962; from then to 2019, per-capita real GDP grew from \$1,918 to \$51,996.² Along these paths of rapid growth, the two countries' wealth distributions evolved differently. In China, wealth concentrated sharply: the top 10% share went from 41% in 1982 to 68% by 2019. In Korea, the distribution barely moved: the top 10% share held near 66% from 1960 to 2019.³ A salient feature of Chinese and Korean development is how deeply the government is involved in it, through its own investment and through redistributive policy, which is why both are often called *developmental states* (Kohli, 2004; Haggard, 2018).

This paper analyzes growth, distribution, and the role of government redistribution in China and Korea. Methodologically, I develop a deep-learning-based Bayesian estimation pipeline for the class of non-stationary heterogeneous-agent growth models that can speak to it, under a bounded-rationality equilibrium framework. The pipeline makes it cheap to trace a whole non-stationary growth path, together with the evolution of the wealth distribution along it, across tens of thousands of simulations, and the fully Bayesian approach carries the uncertainty in the estimated parameters through into the counterfactual experiments.

Based on this computational method, I find the following results, with numerical effects reported at their posterior medians. Total factor productivity (TFP) catch-up is worth a 217% increase in China's per-capita output and 306% in Korea's by 2019, yet it has no clear effect on wealth shares. The education expansion and the fertility decline have positive effects on per-capita output with limited distributional effects. State investment has positive effects on output but may raise wealth concentration: handing households a fifth of state investment reduces the top-10% wealth share by over 23 percentage points in both countries, at a cost of 22% of China's output and 26% of Korea's in 2019. A more generous safety net has negative effects on output and positive effects on consumption at the bottom, although the credible

¹I use "Korea" as shorthand for the Republic of Korea throughout the paper.

²PENN WORLD TABLE (PWT) 11.0 (Feenstra et al., 2015) and World Bank data, constant 2021 US dollars.

³Shares are from the WORLD INEQUALITY DATABASE (Chancel et al., 2022). Korean wealth shares are reported only at decadal benchmarks before 1980, so the 1960 figure is used.

intervals are wide. Among the estimated primitives, idiosyncratic labor-productivity risk has the largest contribution to wealth inequality, worth 6.2 points of China’s top-10% share and 8.6 of Korea’s, ahead of preference heterogeneity.

Two branches of literature explain the East Asian miracle, including that of China and Korea: growth accounting, which decomposes it into factor accumulation and TFP catch-up, and the developmental-state literature, which traces it to the government’s role in steering that accumulation.

Within growth accounting, one strand attributes the miracle mostly to the accumulation of physical and human capital rather than to rising efficiency (Young, 1995, 2003; Krugman, 1994); a second reads the same experience as TFP catch-up toward a moving world frontier (Zhu, 2012; Fernández-Villaverde et al., 2023b); and a third ties growth to the demographic transition, whose changing age structure raises the share of savers and workers (Bloom and Williamson, 1998; Bloom et al., 2003; Zhang et al., 2015), though only within a temporary window (Li et al., 2012). Human capital enters these decompositions in two ways, either as an accumulated factor (Lucas, 1988, 1993) or as the absorptive capacity that governs how fast a follower adopts frontier technology (Nelson and Phelps, 1966; Vandenbussche et al., 2006).

The developmental-state literature instead attributes the miracle to a state that deliberately steers factor allocation toward industrialization and growth (Johnson, 1982; Amsden, 1989; Wade, 1990; Kohli, 2004; Haggard, 2018). Both governments held large state-owned capital stocks during their high-growth periods, Korean state-owned enterprises accounting for over 25% of fixed capital formation by 1980 (Lim, 2003) and Chinese ones for over 50% by 2000,⁴ with state-controlled financial systems channeling credit into infrastructure and strategic industries (Chen et al., 2020; Lane, 2025); Song and Xiong (2024) attribute about half of China’s pre-2008 growth to the career incentives that led local officials to invest in infrastructure. Both also expanded schooling substantially (Barro and Lee, 2013). And both kept public provision minimal during high growth, leaving households to self-insure through saving: China offered essentially no social insurance outside the state-owned sector before 1993 (Hussain, 1994), and Korean social security expenditure stayed near 1% of GNP throughout the 1966–1973 high-growth window (Kohli, 2004).

The growth-accounting literature and the developmental-state literature are thus intertwined: developmental-state policy shapes the very factor accumulation that growth accounting measures. I therefore integrate the growth mechanisms of the developmental states described above into a heterogeneous-agent model. Its primitives are TFP catch-up, the demographic

⁴NBS data, <https://www.stats.gov.cn/sj/ndsj/2001c/f0602c.htm>.

transition, human capital both as an accumulated factor and as the technology absorptive capacity that sets the catch-up speed, state-owned capital, public education spending, and a consumption floor standing in for the safety net. Figure 1 visualizes how they interact. Alongside the idiosyncratic labor-productivity risk of Aiyagari (1994) and the aggregate productivity risk of Krusell and Smith (1998), I model agents’ preferences as random variables,⁵ let human capital accumulate over the working life as in Imai and Keane (2004), and allow hand-to-mouth agents as in Hubbard et al. (1995).




Figure 1: Elements of the Model

Notes: Apricot nodes are exogenous, blue nodes endogenous, as elsewhere in the paper. The dashed arrows are the equilibrium feedback.

This richness comes at a cost. Because the TFP process, demography, and government policies are all indexed to calendar time, the model is non-stationary: the agents’ policy functions change from year to year, and the economy follows a single calendar-time path rather than settling into a time-invariant equilibrium.⁶ Combined with the rich heterogeneity, non-stationarity makes a full rational-expectations equilibrium both empirically implausible and computationally infeasible. It is implausible because, first, households would have to anticipate the joint distribution of future TFP shocks, demographic transitions, and policy reforms over an infinite horizon and, second, they would then have to forecast the entire cross-sectional individual states distribution those shocks induce, since future prices depend on it, a conceptual difficulty for heterogeneous-agent models under rational expectations (Moll, 2026). It is infeasible because

⁵Stachurski and Toda (2019) show that in the standard heterogeneous-agent model with idiosyncratic labor income shocks (Aiyagari, 1994), wealth can never be more heavy-tailed than income.

⁶A stochastic process $(X_t)_{t \in \mathbb{N}}$ is stationary if the distribution of $(X_{t+\Delta t})_{t \in \mathbb{N}}$ is the same as that of $(X_t)_{t \in \mathbb{N}}$ for every Δt (Klenke, 2020, Definition 9.7), the sense in which technology and prices are stationary in the equilibrium of Bewley (1981). In heterogeneous-agent models, “stationary” can mean that the distribution of agents’ states is time-invariant at equilibrium (e.g., Aiyagari, 1994; Gillingham et al., 2022); that the law of motion of the aggregate states is time-invariant (e.g., Krusell and Smith, 1998; Kaplan et al., 2020); or that the simulated economy settles into a time-invariant ergodic set (Judd et al., 2011). All three senses share the property that time is not an argument of the policy function, which fails here.

the curse of dimensionality bites even harder in a non-stationary environment, and the computational methods that exploit stationarity do not apply (e.g., [Auclert et al., 2021](#); [Maliar et al., 2021](#); [Azinovic et al., 2022](#)).

To avoid the computational swamp created by rational expectations, this paper follows the bounded-rationality tradition, which replaces the global rationality of the fully optimizing agent with behavior compatible with the information and computational capacity a decision maker actually possesses ([Simon, 1955](#)). In particular, I build on the *temporary equilibrium* of [Grandmont \(1977\)](#) and [Adam and Marcet \(2011\)](#), in which agents hold subjective beliefs about future prices and choose their actions accordingly. [Cao et al. \(2026\)](#) find that in an automobile-market model a full rational-expectations equilibrium and a naive temporary equilibrium generate very similar outcomes. [Moll \(2026\)](#) surveys these and other departures from rational expectations, and computational progress on making them operational has been rapid.⁷

Building on this tradition, I replace rational expectations with a bounded-rationality concept I call the **certainty-equivalent-anchored (CE-anchored) quasi-rational equilibrium** ([Section 2.7](#)). Each period, agents construct a deterministic forecast of future demography, policies, and prices over a finite horizon, and solve their dynamic optimization problem against that forecast. The forecast is internally consistent: within the forecast window, the predicted wage and interest rate are the *anchor prices* that clear the labor and capital markets in the agents’ own deterministic, noise-free version of the future. Agents then attach subjective price risk around this anchor and optimize against the resulting stochastic prices, so precautionary saving is preserved. As time advances and new shocks arrive, agents observe the realized state and re-forecast, so the economy evolves one period at a time.

[Figure 2](#) summarizes the idea: agents read information off the real economy, plan inside a deterministic imagined economy of their own construction, and act back on the real one. This echoes the “world models” of [Ha and Schmidhuber \(2018\)](#), in which an agent plans inside an internal model of its environment and then acts in the real one. The difference is that the agents’ internal model here is a structurally specified auxiliary economy, not a generative model learned from data.

⁷[Woodford \(2019\)](#) truncates the plan at a finite horizon and closes it with a value function learned from experience; the simulated certainty-equivalent approximation of [Cai and Judd \(2023\)](#) replaces the stochastic future with a single certainty-equivalent path and re-solves as shocks arrive; the N -bounded foresight equilibrium of [Islah and Light \(2025\)](#) traces prices along each possible shock path for a finite number of periods; the self-justified equilibria of [Kubler and Scheidegger \(2025\)](#) restrict agents to forecast a low-dimensional summary of the state; and the structural reinforcement learning of [Yang et al. \(2026\)](#) lets agents track only equilibrium prices, and learn the price dynamics from simulated paths.




Figure 2: Planning in an Imagined World

Notes: The solid loop is what happens in the model economy, the dashed loop what agents merely contemplate. The imagined world is reconstructed each period so agents can forecast the real one.

To make Bayesian estimation of such a non-stationary model with rich heterogeneity feasible, I organize the computational pipeline around amortized optimization (Amos, 2023). Whenever the simulation would otherwise call an expensive solver repeatedly, I clone that solver with a network, in the spirit of approximate dynamic programming (Bertsekas and Tsitsiklis, 1996) and of learning a policy from an expert (Silver et al., 2016). Two devices keep the exact solves that supply the labels affordable: the policy-tree-descendant algorithm (Liu and Iskhakov, 2025) solves each agent’s problem in the imagined deterministic world exactly and grid-free, and a one-step rollout reduces the agents’ stochastic dynamic program to a single look-ahead step (Bertsekas, 2025). The anchor prices are learned differently, by objective learning: they carry no labels, and the network descends the market-clearing loss itself, with a second, small network trained to emulate an exact Gauss–Newton step correcting it where it falls short.

For the intractable likelihood, I estimate the posterior by amortized simulation-based inference, the neural successor to approximate Bayesian computation (Papamakarios and Murray, 2016; Cranmer et al., 2020; Zammit-Mangion et al., 2024): an inference network trained on data simulated from the prior maps an observed series directly to a posterior.⁸

Figure 3 summarizes the pipeline and Section 3 develops it in detail. On a single laptop, generating the training data for these stages takes less than 20 hours, and training the networks takes less than 3 hours (Appendix G).

⁸The classical alternatives are ill-suited here. Bayesian DSGE estimation in Herbst and Schorfheide (2015) requires evaluating the likelihood, which is tractable only after linearizing the model around a stationary steady state, a structure the one-time development transition lacks. The nested fixed-point algorithm (Rust, 1987), a frequentist maximum-likelihood method, needs the same likelihood and so faces the same obstacle. The method of simulated moments (McFadden, 1989; Gouriéroux and Monfort, 1996) matches simulated to empirical moments and needs no likelihood, but with a single realized time series its uncertainty has to come from the bootstrap, which requires many re-estimations, each a search problem nesting numerous simulations.




Figure 3: The Computation Pipeline

Notes: Apricot nodes are the exogenous inputs, blue nodes the stages produced within the pipeline. The prior and the data feed data to the exact optimizers, whose solutions the neural networks clone, save for the anchor prices, which are fit by objective learning from the market-clearing loss; the networks drive the equilibrium simulator; and neural posterior estimation (NPE) maps the simulated data, together with the prior and the observed data, to the posterior over parameters.

Deep learning is the engine that drives this pipeline, for its use in solving heterogeneous-agent models, see [Maliar et al. \(2021\)](#); [Maliar and Maliar \(2022\)](#); [Azinovic et al. \(2022\)](#); [Fernández-Villaverde et al. \(2023a\)](#); [Han et al. \(2026\)](#), surveyed in [Fernández-Villaverde \(2026\)](#). Closest to this paper is [Kase et al. \(2025\)](#), who condition the solution networks on the structural parameters and perform Bayesian inference by training a neural surrogate for the particle-filter likelihood of their model. The major divergence from this literature is twofold. First, I target the non-stationary growth path rather than the stationary equilibrium around a steady state, so there is no ergodic distribution to sample from and the networks must cover the calendar-time transition itself. Second, the pipeline mixes supervised with unsupervised learning: the decision rules are cloned from exact-solver labels, while the anchor prices are learned unsupervised from the market-clearing loss.

The rest of the paper is organized as follows. [Section 2](#) develops the structural model and defines the *CE-anchored quasi-rational equilibrium*. [Section 3](#) sets out the computational pipeline and reports the estimates. [Section 4](#) turns to the counterfactual experiments, which trace the mechanisms behind growth and (re)distribution, and [Section 5](#) concludes.

2 Model

This section presents the model, which is built upon the concept of *certainty-equivalent-anchored quasi-rational equilibrium* ([Section 2.7](#)) and structured as a directed probabilistic model, i.e., a Bayes net. A Bayes net is a joint probability distribution over a collection of random variables, represented by nodes of a directed acyclic graph (DAG). Every node is generated by exactly one structural equation, and the variables entering that equation are its parents. The graph

is therefore a structural causal model (Pearl, 2009), which gives meaning to an intervention, or what structural econometrics calls a counterfactual experiment: it replaces the structural equation of one node while leaving every other equation intact, so the change propagates to exactly the distributions of all that node’s descendants. See Pearl et al. (2016) and Murphy (2023, Chapter 4) for textbook treatments of causal graphs and Bayes nets.

The model is composed of several interconnected components, each governing the evolution of a part of the economy. For each component, I present the structural equations alongside a DAG that displays the causal channels.

The prior distributions for the model’s parameters are provided in Table 1, with detailed rationales available in Appendix B. While the priors are separated to streamline the presentation of the model’s core architecture, they are an integral part of the Bayes net since they reflect “the initial plausibility of each conjectured cause of the data” (McElreath, 2020).

In the following, for notational brevity, I define $X_{t_1:t_2} := [X_{t_1}, X_{t_1+1}, \dots, X_{t_2}]$ for any stochastic process or real sequence $\{X_t\}_{t \in \mathbb{N}}$ and $t_1 \leq t_2$, allowing $t_2 = \infty$ with the convention $X_{t_1:\infty} := [X_{t_1}, X_{t_1+1}, \dots]$. I also write $\hat{X}_{t \rightarrow t_1:t_2}$ for the prediction of $X_{t_1:t_2}$ made at time t .

2.1 Demography, Family and Agents

I model China and Korea as two economies populated by dynastic families. Let $j \in \mathcal{C}$ index country, where $\mathcal{C} := \{cn, kr\}$, with cn denoting China and kr denoting Korea. Time is discrete and annual. Country j starts its catch-up process at $t_{s,j}$. In this paper, $t_{s,kr} = 1962$ and $t_{s,cn} = 1982$.

Every agent in the model belongs to a dynastic family: its members share preference parameters (Section 2.4) and are linked by bequests. For each country j , let $(I, \mathcal{A}, \mathbf{P}_j)$ be the probability space of families. Here, $I := [0, 1]$ is the common family-type space, $\mathcal{A}$ is the Borel σ -algebra on I , and $\mathbf{P}_j$ is the country-specific probability measure on $\mathcal{A}$. I use the tuple $(i, j) \in I \times \mathcal{C}$ to denote the family with type i living in country j .

Agents enter the model at adjusted age $a = 1$, which corresponds to 15 years old in the data. An entrant in country j with entry year t_0 studies or works from age $a = 1$ in calendar year t_0 , then dies at the end of age $a = d_{j,t_0}$, in calendar year $t_0 + d_{j,t_0} - 1$. The variable d_{j,t_0} denotes the life expectancy, measured in adjusted-age years, of the cohort entering country j at t_0 . Let $a_{j,max,t} := \max\{a \in \mathbb{N} : d_{j,t-a+1} \geq a\}$ denote the age of the oldest cohort still alive in country j at t , so the cohorts aged 1 through $a_{j,max,t}$ are the ones alive at t . Hereafter, I use “age” as shorthand for “adjusted age”.

At each time t , each family (i, j) has one entrant of age $a = 1$. Let $N_{j,t}$ denote the size of the age-1 entrant cohort in country j at time t , standardized so that each country's base-year population equals one. Each family therefore holds one agent of every age $a = 1, \dots, a_{j,max,t}$, the cohort aged a having entered at $t - a + 1$ with size $N_{j,t-a+1}$. In the following, I use the normalized notation $\iota := (i, j, t_0)$ for the agent of family type i who enters country j at time t_0 .⁹ The notations ι_i , ι_j , and ι_{t_0} denote her family type, country, and entry year.

At the start of her entry year, each agent receives a bequest from her family's recent exits, as formalized in Section 2.3. Upon entry, agent ι is assigned an education level $e_\iota \in \{e_b, e_h, e_c\}$: e_b (*below high school*), e_h (*high school*), and e_c (*college*). The assignment is determined by her family index and the government's education policy, as introduced in Section 2.3. The education level determines each agent's labor supply process as introduced in Section 2.4. Let a_e denote the age at which agents with education level e begin working: $a_{e_b} = 1$ for agents with e_b , who work from age $a = 1$; $a_{e_h} = 4$ for agents with e_h , who study from $a = 1$ to $a = 3$; and $a_{e_c} = 8$ for agents with e_c , who study from $a = 1$ to $a = 7$.

An agent $\iota = (i, j, t_0)$ exits the model after completing age d_{j,t_0} in calendar year $t_0 + d_{j,t_0} - 1$. At that time, she leaves a bequest—her resources remaining after final-age consumption—in equal shares to the next three entrant cohorts of the same family. Because the life expectancy d_{j,t_0} varies across cohorts, a family need not have a member exiting in every calendar year. Spreading each bequest over three entrant cohorts ensures that every entrant still receives a bequest.

The entrant-cohort size path, $\{N_{j,t} : j \in \mathcal{C}, t \in \mathbb{N}\}$, and the life-expectancy path, $\{d_{j,t} : j \in \mathcal{C}, t \in \mathbb{N}\}$, are exogenous in this model.

2.2 Production

The production side of each economy is represented by a competitive firm that produces a single final good with capital and labor. I use a CES technology so that factor income shares need not be fixed: as the economy accumulates capital and its demographic structure changes, the relative importance of capital and labor in production can also change. For country j in year t , output is

$$Y_{j,t} = A_{j,t} \left(\alpha_{Y,j} K_{j,t}^{\rho_{Y,j}} + (1 - \alpha_{Y,j}) L_{j,t}^{\rho_{Y,j}} \right)^{1/\rho_{Y,j}}, \quad \rho_{Y,j} = \frac{\sigma_{Y,j} - 1}{\sigma_{Y,j}}, \quad (1)$$

where $A_{j,t}$ is total factor productivity (TFP), $K_{j,t}$ is aggregate capital, and $L_{j,t}$ is the aggregate labor input. The parameter $\alpha_{Y,j}$ is the CES share parameter, and $\sigma_{Y,j}$ is the elasticity of

⁹For a fixed country-year (j, t_0) , the population-scaled label $(iN_{j,t_0}, j, t_0) \in [0, N_{j,t_0}] \times \mathcal{C} \times \mathbb{N}$ and the normalized label (i, j, t_0) identify the same entrant. Since N_{j,t_0} is fixed for the country-year, the two labels are equivalent.

substitution between capital and labor.

I next describe the construction of these aggregate inputs, starting with TFP.

Total Factor Productivity $A_{j,t}$ I model TFP along the lines of [Fernández-Villaverde et al. \(2023b\)](#), in which countries learn from a common technology frontier and approach a country-specific ceiling through catch-up, but with two departures: I endogenize the catch-up speed through an education-weighted density of human capital, which places more weight on the human capital of more educated workers, and I add a stochastic TFP shock.

I assume that all countries can learn from a common world technology frontier. Let $A_{f,t}$ denote the frontier TFP level in year t . The frontier grows at the constant rate g_f :

$$A_{f,t+1} = (1 + g_f)A_{f,t}. \quad (2)$$

The country-specific TFP is the product of the frontier, a country-specific ceiling, the country's position relative to that ceiling, and a TFP shock:

$$A_{j,t} = \theta_{A,j} \lambda_{j,t} A_{f,t} \exp(U_{j,t}), \quad (3)$$

where $\theta_{A,j}$ is country j 's technology ceiling, measured as the asymptotic fraction of frontier TFP that the country can attain, reflecting institutional or structural limits. The term $\lambda_{j,t}$ is the fraction of this ceiling that the country has already reached, and $U_{j,t}$ is a TFP shock.

The ratio $\lambda_{j,t}$ evolves through catch-up:

$$\log \lambda_{j,t} = (1 - v_{j,t}) \log \lambda_{j,t-1}. \quad (4)$$

Since $v_{j,t} \in [0, 1]$, the recursion contracts $\lambda_{j,t}$ toward 1. The term $v_{j,t}$ is therefore the adoption speed in period t : given the previous period's position $\lambda_{j,t-1}$, a higher $v_{j,t}$ moves the country more quickly toward its ceiling in period t . This formulation also captures the low-hanging-fruit feature of technology diffusion: early steps on the technology ladder generate larger proportional gains, while further movement toward the ceiling becomes progressively more difficult.

The main departure from [Fernández-Villaverde et al. \(2023b\)](#) is that the catch-up speed $v_{j,t}$ is endogenous and rises with the workforce's human capital. Following the literature in which human capital governs the speed of technology adoption ([Nelson and Phelps, 1966](#); [Benhabib and Spiegel, 1994](#); [Barro, 2013](#)), I tie $v_{j,t}$ to an education-weighted density $\mathcal{H}_{j,t} \in [0, 1]$

of the human capital of the cohorts aged $a \geq a_{ec}$,

$$v_{j,t} = v_{max} \mathcal{H}_{j,t}, \quad \mathcal{H}_{j,t} = \frac{\sum_{a=a_{ec}}^{a_{j,max,t}} N_{j,t-a+1} \int \omega_{e(i,j,t-a+1)} h_{(i,j,t-a+1),a} \mathbf{P}_j(di)}{\omega_c \sum_{a=a_{ec}}^{a_{j,max,t}} N_{j,t-a+1} h_{max,a}}, \quad (5)$$

where v_{max} is the maximum adoption speed. The education weights $\omega_{e_i} \in \{\omega_b, \omega_h, \omega_c\}$ satisfy $\omega_c \geq \omega_h \geq \omega_b \equiv 1$, so more-educated workers receive greater weight, consistent with their disproportionate role in absorbing frontier technologies and generating innovations (Vandenbussche et al., 2006). The denominator normalizes $\mathcal{H}_{j,t}$ within $[0, 1]$ by the maximal attainable human capital $h_{max,a}$, that of a college-educated agent who works full-time at every age, formally defined in Section 2.4.

Let $\mathbf{h}_{j,t}$ denote the cross-sectional distribution of human capital and $\mathbf{N}_{j,t} := N_{j,t-a_{j,max,t}+1:t}$ the vector of alive cohort sizes. The cohort sizes $\mathbf{N}_{j,t}$ and the human-capital distribution $\mathbf{h}_{j,t}$ together determine the education-weighted density $\mathcal{H}_{j,t}$ over the cohorts aged $a \geq a_{ec}$, which in turn sets the catch-up speed and the evolution of the TFP-to-ceiling ratio. Figure 4 summarizes this technology catch-up mechanism as a computational graph. Apricot nodes are exogenous random variables and blue nodes are endogenous ones, here and in the DAGs that follow.




Figure 4: Technology Catch-up Process

Notes: $\mathbf{N}_{j,t}$ is the vector of alive cohort sizes; $\mathbf{h}_{j,t}$ is the cross-sectional distribution of human capital; $\mathcal{H}_{j,t}$ is the education-weighted human-capital density; $v_{j,t}$ is the catch-up speed; and $\lambda_{j,t}$ is the TFP-to-ceiling ratio.

The TFP shock captures transitory country-specific deviations from the frontier-and-catch-up component of productivity. I model $U_{j,t}$ as an AR(1) process whose innovation mean is set so that the multiplicative shock $\exp(U_{j,t})$ has unit mean in the stationary distribution:

$$U_{j,t} = \eta_{j,u} U_{j,t-1} + \epsilon_{u,j,t}, \quad \epsilon_{u,j,t} \sim \mathcal{N}\left(-\frac{\sigma_{j,u}^2}{2(1 + \eta_{j,u})}, \sigma_{j,u}^2\right). \quad (6)$$

The shock is initialized at the catch-up start year $t_{s,j}$ from the stationary distribution of (6), so that, together with the negative innovation mean, the multiplicative shock has unit mean in

every year, $\mathbb{E}[\exp(U_{j,t})] = 1$. Figure 5 summarizes this shock process.




Figure 5: TFP Shock Process

Notes: $\epsilon_{u,j,t}$ is the TFP-shock innovation; and $U_{j,t}$ is the TFP shock.

Aggregate Capital Stock $K_{j,t}$ The capital input in the production function is effective aggregate capital. It combines private capital, $K_{p,j,t}$, and state capital, $K_{s,j,t}$, through a CES aggregator:

$$K_{j,t} = \left(\alpha_{K,j} K_{p,j,t}^{\rho_{K,j}} + (1 - \alpha_{K,j}) K_{s,j,t}^{\rho_{K,j}} \right)^{1/\rho_{K,j}}, \quad \rho_{K,j} = \frac{\sigma_{K,j} - 1}{\sigma_{K,j}}. \quad (7)$$

The parameter $\alpha_{K,j}$ is the distribution parameter, and $\sigma_{K,j}$ is the elasticity of substitution between private and state capital. I impose $\sigma_{K,j} < 1$, so the two forms of capital are gross complements. This assumption reflects the role of state capital as public infrastructure and productive capacity that raises the usefulness of private investment.

The private-capital component is the economy-wide sum of individual capital holdings across all alive cohorts:

$$K_{p,j,t} = \sum_{a=1}^{a_{j,max,t}} N_{j,t-a+1} \int k_{(i,j,t-a+1),a} \mathbf{P}_j(di). \quad (8)$$

I denote the cross-sectional distribution of alive agents' capital at t by $\mathbf{k}_{j,t}$. Together with the alive cohort-size vector $\mathbf{N}_{j,t}$, this distribution determines aggregate private capital, $K_{p,j,t}$. The formation of state capital $K_{s,j,t}$ through government investment is introduced in Section 2.3. Aggregate private capital and state capital then combine through (7) to form the effective aggregate capital input, $K_{j,t}$. Figure 6 summarizes these dependencies.




Figure 6: Aggregate Capital Formation

Notes: $\mathbf{k}_{j,t}$ is the cross-sectional distribution of individual capital holdings; $\mathbf{N}_{j,t}$ is the vector of alive cohort sizes; $K_{p,j,t}$ is aggregate private capital; $K_{s,j,t}$ is state capital; and $K_{j,t}$ is effective aggregate capital.

Aggregate Labor Supply $L_{j,t}$ Labor input is efficient labor, measured as human capital in use and scaled by an idiosyncratic labor-productivity shock. For an agent i at age a , this efficient labor is the product of human capital $h_{i,a}$, working time $\ell_{i,a} \in [0, 1]$, and the shock $\exp(\varepsilon_{i,a})$, where $\varepsilon_{i,a}$ follows the AR(1) law with $\mathbb{E}[\exp(\varepsilon_{i,a})] = 1$. Since working time is chosen conditional on the realized shock, the aggregate labor supply depends on the joint distribution of human capital, working time, and productivity shocks $(\mathbf{h}, \mathbf{\ell}, \mathbf{\varepsilon})_{j,t}$, together with the alive cohort-size vector $\mathbf{N}_{j,t}$:

$$L_{j,t} = \sum_{a=1}^{a_{j,max,t}} N_{j,t-a+1} \int h_{(i,j,t-a+1),a} \ell_{(i,j,t-a+1),a} \exp(\varepsilon_{(i,j,t-a+1),a}) \mathbf{P}_j(di). \quad (9)$$

Figure 7 summarizes this computation.




Figure 7: Aggregate Labor Input

Notes: $(\mathbf{h}, \mathbf{\ell}, \mathbf{\varepsilon})_{j,t}$ is the joint distribution of human capital, working time, and productivity shocks; $\mathbf{N}_{j,t}$ is the vector of alive cohort sizes; and $L_{j,t}$ is the aggregate labor input.

Wage and Interest Rates The nested CES structure lets marginal products be written in terms of two sets of input income shares. Let $s_{p,j,t}$ be the private-capital share inside the aggregate-capital bundle and let $s_{s,j,t}$ be its state-capital counterpart. Similarly, let $s_{K,j,t}$ be the aggregate-capital share in final-good production and let $s_{L,j,t}$ be the labor share, defined as

$$\begin{aligned} s_{p,j,t} &= \frac{\alpha_{K,j} K_{p,j,t}^{\rho_{K,j}}}{\alpha_{K,j} K_{p,j,t}^{\rho_{K,j}} + (1 - \alpha_{K,j}) K_{s,j,t}^{\rho_{K,j}}}, \\ s_{s,j,t} &= 1 - s_{p,j,t}, \\ s_{K,j,t} &= \frac{\alpha_{Y,j} K_{j,t}^{\rho_{Y,j}}}{\alpha_{Y,j} K_{j,t}^{\rho_{Y,j}} + (1 - \alpha_{Y,j}) L_{j,t}^{\rho_{Y,j}}}, \\ s_{L,j,t} &= 1 - s_{K,j,t}. \end{aligned} \quad (10)$$

Under perfect competition, firms pay factors their marginal products. I use a common depreciation rate, δ , for both private and state capital, so the wage and the capital returns net

of depreciation are

$$\begin{aligned}
w_{j,t} &= \frac{Y_{j,t}}{L_{j,t}} s_{L,j,t}, \\
r_{p,j,t} &= \frac{Y_{j,t}}{K_{p,j,t}} s_{K,j,t} s_{p,j,t} - \delta, \\
r_{s,j,t} &= \frac{Y_{j,t}}{K_{s,j,t}} s_{K,j,t} s_{s,j,t} - \delta,
\end{aligned} \tag{11}$$

where $w_{j,t}$ is the wage for one unit of efficient labor, $r_{p,j,t}$ is the private-capital return, $r_{s,j,t}$ is the state-capital return. State capital therefore affects households through wages, taxes, public revenues, and equilibrium prices. [Figure 8](#) gives a compact overview of the production and factor-price computations.




Figure 8: Production and Factor Prices

Notes: $A_{f,t}$ is frontier TFP; $\lambda_{j,t}$ is the TFP-to-ceiling ratio; $U_{j,t}$ is the TFP shock; $A_{j,t}$ is TFP; $K_{p,j,t}$ and $K_{s,j,t}$ are private and state capital; $L_{j,t}$ is aggregate labor; $Y_{j,t}$ is output; and the factor-price node collects $w_{j,t}$, $r_{p,j,t}$, and $r_{s,j,t}$.

2.3 Government

The government collects tax revenue and invests while providing social insurance and education for the agents. The multiple roles give the government significant power to not only stimulate the economic growth but also shape the wealth distribution. This subsection introduces the government's expenditure-revenue structure.

The government's expenditure in each country j includes three components:

- (i) **Social welfare $W_{j,t}$ that provides a consumption floor:** Following [Hubbard et al. \(1995\)](#), I model social welfare as a consumption floor $\underline{x}_{j,t}$ provided to all agents with insufficient cash-on-hand. Let $\mathcal{T}_{l,j,t}(\cdot)$ denote the labor income tax liability function in country j at time t , which captures the progressive tax schedule. Writing agent ι 's gross labor income at age a as $w_{j,t} h_{\iota,a} \ell_{\iota,a} \exp(\varepsilon_{\iota,a})$, her after-tax labor income is

$$y_{\iota,a} = w_{j,t} h_{\iota,a} \ell_{\iota,a} \exp(\varepsilon_{\iota,a}) - \mathcal{T}_{l,j,t}(w_{j,t} h_{\iota,a} \ell_{\iota,a} \exp(\varepsilon_{\iota,a})). \tag{12}$$

Her pre-floor cash-on-hand $z_{\iota,a}$ is her last-period saving $k_{\iota,a}$ scaled by the after-tax return, plus this labor income:

$$z_{\iota,a} = (1 + r_{p,j,t} - \max\{0, \tau_{k,j,t} r_{p,j,t}\}) k_{\iota,a} + y_{\iota,a}.$$

The consumption floor tops up any shortfall in pre-floor cash-on-hand, giving the post-floor cash-on-hand $x_{\iota,a} = \max\{z_{\iota,a}, \underline{x}_{j,t}\}$ of (39). Thus, the total social welfare expenditure is:

$$W_{j,t} = \sum_{a=1}^{a_{j,max,t}} N_{j,t-a+1} \int \max\{\underline{x}_{j,t} - z_{(i,j,t-a+1),a}, 0\} \mathbf{P}_j(di). \quad (13)$$

Since the consumption floor serves as a simplified representation of complex real-world social welfare systems, I model it as a stochastic process that reflects the evolution of welfare policies.¹⁰ To capture potential structural changes in government welfare policies while avoiding parameter proliferation, I introduce a break point $t_{g,j}$ for each country j . Based on major political-economic structural change, I set $t_{g,kr} = 1988$ and $t_{g,cn} = 2012$.¹¹ The consumption floor is a piecewise-constant fraction of the economy-wide wage level:

$$\underline{x}_{j,t} = \begin{cases} \Lambda_{j,1} w_{j,t} & \text{if } t \leq t_{g,j}, \\ \Lambda_{j,2} w_{j,t} & \text{if } t > t_{g,j}. \end{cases} \quad (14)$$

$\Lambda_{j,1}$ and $\Lambda_{j,2}$ are the ratios of the consumption floor to the economy-wide wage level, $w_{j,t}$, for the two regimes before and after the structural break. The wage $w_{j,t}$ is the income for one unit of efficient labor in (11).

Computing the total welfare expenditure in (13) requires identifying all agents whose cash-on-hand is at the consumption floor. This necessitates knowing the entire cross-sectional distribution of cash-on-hand at time t , which is denoted by

$$\mathbf{x}_{j,t} := \left\{ x_{(i,j,t-a+1),a} : i \in I, 1 \leq a \leq a_{j,max,t} \right\}.$$

Conceptually, $\mathbf{x}_{j,t}$ is an infinite-dimensional object, posing a significant computational

¹⁰This time-varying specification contrasts with the structural labor economics literature, where the consumption floor is typically modeled as a time-invariant constant (De Nardi et al., 2010; French and Jones, 2011; Fan et al., 2024) for estimation or calibration.

¹¹The Sixth Republic of Korea was established in February 1988, marking a transition to democratic governance, and pension programs were introduced for the general public in the same year. See Koh (2015) for the evolution of Korea's social welfare system. The 18th National Congress of the Chinese Communist Party was held in November 2012, establishing new leadership and policy directions. This National Congress preceded the Targeted Poverty Alleviation policy, proposed in 2013 and officially adopted in 2014, which aimed to eliminate poverty by 2020. See Zhou et al. (2023); Li et al. (2025) for details on Targeted Poverty Alleviation policies.

challenge that I address in [Sections 2.5 to 2.7](#). The dependencies for computing the total welfare expenditure, $W_{j,t}$, are summarized by the DAG in [Figure 9](#).




Figure 9: Welfare Expense

Notes: $w_{j,t}$: wage; $x_{j,t}$: consumption floor; $x_{j,t}$: cross-sectional cash-on-hand distribution; $N_{j,t-a_{j,max},t+1:t}$: alive cohort sizes; $W_{j,t}$: welfare expenditure.

- (ii) **Education Investment $H_{j,t}$ that builds human capital:** For the entrant cohort at time t in each country j , the government sets enrollment ratios $\epsilon_{h,j,t}$ (high school only) and $\epsilon_{c,j,t}$ (college). These ratios satisfy $0 \leq \epsilon_{c,j,t}$, $0 \leq \epsilon_{h,j,t}$, and $\epsilon_{c,j,t} + \epsilon_{h,j,t} \leq 1$. The government then assigns education levels to each entrant $\iota := (i, j, t)$ based on her study ability $\phi_\iota := \phi(\iota_i)$. Without loss of generality, ϕ is assumed to be an $\mathcal{A}$ -measurable function such that, under $\mathbf{P}_j$, $\phi(i)$ is uniformly distributed on $(0, 1)$ in each country j , since only the ranking of study ability matters:

$$e_\iota = \begin{cases} e_c & \text{if } \mathbf{P}_j[\{i' \in I : \phi(i') \geq \phi_\iota\}] \leq \epsilon_{c,j,t}, \\ e_h & \text{if } \epsilon_{c,j,t} < \mathbf{P}_j[\{i' \in I : \phi(i') \geq \phi_\iota\}] \leq \epsilon_{h,j,t} + \epsilon_{c,j,t}, \\ e_b & \text{otherwise.} \end{cases} \quad (15)$$

With the students-to-teachers ratio θ_κ , the total education investment is:

$$H_{j,t} = \frac{w_{j,t}}{\theta_\kappa} \left[\sum_{t'=t-a_{e_h}+2}^t \epsilon_{h,j,t'} N_{j,t'} + \sum_{t'=t-a_{e_c}+2}^t \epsilon_{c,j,t'} N_{j,t'} \right], \quad (16)$$

where the per-student education cost is $w_{j,t}/\theta_\kappa$: one teacher's wage $w_{j,t}$ shared among θ_κ students, and each sum runs over the entrant cohorts still in school at t . The computational graph of education investment is shown in [Figure 10](#).




Figure 10: Education Investment

Notes: $\epsilon_{h,j,t-2:t}$: high school enrollment rates; $\epsilon_{c,j,t-6:t}$: college enrollment rates; $w_{j,t}$: wage; $N_{j,t-6:t}$: entrant cohort sizes; $H_{j,t}$: education expenditure.

- (iii) **State-owned capital accumulation $K_{s,j,t+1}$ that accumulates state-owned capital:** The government carries fiscal resources into the next period as state-owned capital, $K_{s,j,t+1}$. Thus, the choice of $K_{s,j,t+1}$ is the government's state-capital accumulation decision at time t . The subscript "s" denotes state ownership, distinguishing it from aggregate private capital $K_{p,j,t}$ defined in (8).

This accumulation decision, $K_{s,j,t+1}$, links the government's budget across time. As it is determined by the balance of expenditures and revenues, I defer showing its computational graph (Figure 14) until after the revenue side is detailed.

The government's revenue in each country j consists of four sources:

- (i) **Tax revenue $T_{k,j,t}$ from capital income:** The government taxes capital income earned by agents. Let $k_{\iota,a}$ denote the capital held by agent ι at age a , $r_{p,j,t}$ denote the pre-tax private-capital return in country j at time t , and $\tau_{k,j,t}$ denote the tax rate in country j on capital income. The aggregate private capital stock $K_{p,j,t}$ is defined in (8).

To ensure that the government does not subsidize agents when private-capital returns are negative, the total tax revenue from capital income is calculated as:

$$T_{k,j,t} = \max\{0, \tau_{k,j,t} r_{p,j,t} K_{p,j,t}\}. \quad (17)$$

Given the complexity of capital income tax schemes and the fact that both countries employ relatively low tax rates to stimulate saving and investment, I model the effective capital-income tax rate as piecewise constant in time: China's rate is a single constant, $\tau_{k,cn,t} \equiv \tau_{k,cn}$, while Korea's breaks at the 1997 Asian financial crisis, equal to $\tau_{k,kr,1}$ for $t \leq 1997$ and $\tau_{k,kr,2}$ for $t > 1997$, reflecting the higher effective capital taxation of its more mature, post-crisis tax regime.

As (8) and (17) show, the calculation of capital income tax revenue, $T_{k,j,t}$, depends on the aggregate private capital stock, $K_{p,j,t}$. This aggregate is itself computed by summing over the entire cross-sectional distribution of individual capital holdings, which I denote as:

$$\mathbf{k}_{j,t} := \left\{ k_{(i,j,t-a+1),a} : i \in I, 1 \leq a \leq a_{j,max,t} \right\}.$$

Similar to the cash-on-hand distribution $\mathbf{x}_{j,t}$, this is an infinite-dimensional object that poses a significant computational challenge. The DAG in Figure 11 illustrates these dependencies.




Figure 11: Tax Revenue from Capital Income

Notes: $r_{p,j,t}$: private-capital return; $k_{j,t}$: private capital distribution; $N_{j,t-a_{j,max,t}+1:t}$: alive cohort sizes; $K_{p,j,t}$: private capital; $T_{k,j,t}$: capital tax revenue.

- (ii) **Tax revenue $T_{l,j,t}$ from labor income:** The government also taxes labor income earned by agents, through the liability function $\mathcal{T}_{l,j,t}(\cdot)$ of (12). The aggregate labor input $L_{j,t}$ at time t is the $h\ell \exp(\varepsilon)$ aggregate across all working-age agents, computed through (9). The tax base below is the individual labor income and therefore carries the same idiosyncratic shock $\exp(\varepsilon_{\iota,a})$.

The total tax revenue from labor income is then:

$$T_{l,j,t} = \sum_{a=1}^{a_{j,max,t}} N_{j,t-a+1} \int \mathcal{T}_{l,j,t}(w_{j,t} h_{(i,j,t-a+1),a} \ell_{(i,j,t-a+1),a} \exp(\varepsilon_{(i,j,t-a+1),a})) \mathbf{P}_j(di). \quad (18)$$

The computation in (18) depends on the cross-sectional distribution of individual labor incomes, as shown in Figure 12.




Figure 12: Tax Revenue from Labor Income

Notes: $w_{j,t}$: wage; $(h, l, \epsilon)_{j,t}$: human-capital, working-time, and productivity-shock distribution; $N_{j,t-a_{j,max,t}+1:t}$: alive cohort sizes; $T_{l,j,t}$: labor tax revenue.

- (iii) **Additional tax revenue $T_{m,j,t}$ from transfer:** The government at t can also tax the transfer received by new entrants. Let $\mathbf{g}_{\iota}$ denote the bequest left by agent ι upon exit, and let $b_{(i,j,t)}$ denote the bequest received by the entrant of family type i in country j at time t . For each calendar year t , define the set of same-family agents who exit at the end of t as

$$\mathcal{D}_{i,j,t} := \{(i, j, t_0) : t_0 + d_{j,t_0} - 1 = t\}.$$

With the three-year bequest kernel,

$$b_{(i,j,t)} = \sum_{q=1}^3 \left[\left(\prod_{s=t-q+1}^t (1 + r_{p,j,s} - \max\{0, \tau_{k,j,s} r_{p,j,s}\}) \right) \sum_{\iota \in \mathcal{D}_{i,j,t-q}} \frac{N_{j,\iota_{t_0}}}{N_{j,t}} \frac{\mathbf{g}_\iota}{3} \right]. \quad (19)$$

Thus, the received bequest combines one-third shares of bequests left by same-family agents who exited at the ends of $t - 1$, $t - 2$, and $t - 3$, with each share adjusted by the relative size of the exiting cohort and the entrant cohort. Since entrants receive bequests as beginning-of-period resources at t , a bequest left at the end of $t - q$ accumulates after-tax private-capital returns from $t - q + 1$ through t . Let $\tau_{m,j,t}$ denote the transfer tax rate. The total tax revenue from transfer is:

$$T_{m,j,t} = \tau_{m,j,t} N_{j,t} \int b_{(i,j,t)} \mathbf{P}_j(di). \quad (20)$$

The taxable amount is therefore the sum of bequests received by all entrants at t .

Since the transfer in the model mirrors bequest in the real world, and given that China has no inheritance tax while Korea has maintained an inheritance tax for a long time,¹² $\tau_{m,j,t}$ is modeled as:

$$\begin{aligned} \tau_{m,cn,t} &= 0 \text{ for all } t, \\ \tau_{m,kr,t} &= \tau_{m,kr} \text{ for all } t. \end{aligned} \quad (21)$$

The calculation of transfer tax revenue in (20) requires the cross-sectional distribution of bequests received by entrants, $\mathbf{b}_{j,t} := \{b_{(i,j,t)}\}_{i \in I}$. Computing $\mathbf{b}_{j,t}$ in turn requires the cross-sectional distributions of terminal bequests left by same-family agents who exited at the ends of $t - 1$, $t - 2$, and $t - 3$, $\mathbf{g}_{j,t-3:t-1} := \{\mathbf{g}_\iota : \iota \in \mathcal{D}_{i,j,t-3:t-1}\}$, together with the private-capital return path. Similar to the cash-on-hand and capital distributions, these objects are infinite-dimensional and pose a significant computational challenge. The computational graph for this process is shown in Figure 13.




Figure 13: Tax Revenue from Transfer

Notes: $\mathbf{g}_{j,t-3:t-1}$: terminal-bequest distributions (exit years $t - 1, t - 2, t - 3$); $r_{p,j,t-2:t}$: private-capital return path; $\mathbf{b}_{j,t}$: received-bequest distribution; $N_{j,t}$: entrant cohort size; $T_{m,j,t}$: transfer tax revenue.

¹²See OECD report, *Inheritance Taxation in OECD Countries*: <https://doi.org/10.1787/e2879a7d-en>.

- (iv) **Investment revenue** $(1 + r_{s,j,t})K_{s,j,t}$ **from state-owned capital:** The government earns investment returns on its capital holdings at the state-capital return.

Given the expenditure-revenue structure above, the government's budget constraint at time t is:

$$\begin{aligned} W_{j,t} + H_{j,t} + K_{s,j,t+1} &= \zeta_{j,t} (T_{k,j,t} + T_{l,j,t} + T_{m,j,t} + (1 + r_{s,j,t})K_{s,j,t}), \\ W_{j,t} + H_{j,t} &\leq \zeta_{j,t} (T_{k,j,t} + T_{l,j,t} + T_{m,j,t} + (1 + r_{s,j,t})K_{s,j,t}), \end{aligned} \quad (22)$$

where $\zeta_{j,t} \in (0, 1]$ reflects government fiscal efficiency in country j . The second line imposes fiscal feasibility and ensures $K_{s,j,t+1} \geq 0$.

To capture governance changes following major political-economic transitions, I model $\zeta_{j,t}$ as a piecewise constant function with breaks at $t_{g,j}$. The democratic transition in Korea and China's anti-corruption campaign¹³ are expected to impact government fiscal efficiency:

$$\zeta_{j,t} = \begin{cases} \zeta_{j,1} & \text{if } t \leq t_{g,j}, \\ \zeta_{j,2} & \text{if } t > t_{g,j}. \end{cases} \quad (23)$$

where $\zeta_{j,1}$ and $\zeta_{j,2}$ are the fiscal efficiency parameters before and after the structural break, respectively. The computational graph for state-owned capital accumulation is shown in Figure 14 below.




Figure 14: State-Owned Capital Accumulation

Notes: $T_{k,j,t}$: capital tax revenue; $T_{l,j,t}$: labor tax revenue; $T_{m,j,t}$: transfer tax revenue; $r_{s,j,t}$: state-capital return; $W_{j,t}$: welfare expenditure; $H_{j,t}$: education expenditure; $K_{s,j,t}$: state-owned capital.

This completes the specification of the government sector. The computational graphs have revealed a recurring theme: key fiscal variables, such as welfare expenditure and tax revenues, depend on the endogenous cross-sectional distributions of individual cash-on-hand ($\mathbf{x}_{j,t}$), capital holdings ($\mathbf{k}_{j,t}$), and received bequests ($\mathbf{b}_{j,t}$). Since these distributions arise from the collective decisions of individual agents, characterizing the agent's decision problem is the necessary next step to make the government model computable. I now turn to this task.

¹³The Chinese government launched a sustained large-scale anti-corruption campaign following the 18th National Congress of the Chinese Communist Party in 2012. See Fang and Zhang (2025) for details.

2.4 Agents' Preferences, Labor Supply, and Human Capital Accumulation

This subsection specifies agents' preferences, labor supply, and human-capital accumulation.

Preferences Let $\iota = (i, j, t_0)$ denote an agent of family type i , country j , and entry year t_0 . At age a , the agent's period utility from consumption and working time is

$$u_\iota(c, \ell) = \frac{c^{1-\gamma_i} - 1}{1 - \gamma_i} - \chi_i \frac{\ell^{1+\nu_i}}{1 + \nu_i}, \quad (24)$$

for $c > 0$ and $\ell \in [0, 1]$. The parameter γ_i is the coefficient of relative risk aversion, with the logarithmic limit used when $\gamma_i = 1$; χ_i governs the weight on the disutility of work; and $\nu_i > 0$ determines the convexity of the disutility of work.

At the exit age $a = d_{j,t_0}$, the agent also receives terminal bequest utility

$$u_\iota^{\text{beq}}(\mathbf{g}) = \psi_i \frac{\mathbf{g}^{1-\gamma_i} - 1}{1 - \gamma_i}, \quad (25)$$

for bequest amount $\mathbf{g} > 0$, again interpreted as the logarithmic limit when $\gamma_i = 1$. The parameter $\psi_i \geq 0$ is the relative weight on the bequest motive.

Labor Supply and Human Capital Accumulation Agent ι supplies labor $q_{\iota,a}$ in efficiency units according to

$$q_{\iota,a} = h_{\iota,a} \ell_{\iota,a} \exp(\varepsilon_{\iota,a}), \quad (26)$$

where $h_{\iota,a}$ is human capital, $\ell_{\iota,a} \in [0, 1]$ is working time, and $\varepsilon_{\iota,a}$ is an idiosyncratic labor productivity shock. The shock follows the autoregressive process

$$\varepsilon_{\iota,a} = \eta_\varepsilon \varepsilon_{\iota,a-1} + \epsilon_{\varepsilon,\iota,a}. \quad (27)$$

Here η_ε is the persistence parameter and $\epsilon_{\varepsilon,\iota,a}$ is the age- a innovation,

$$\epsilon_{\varepsilon,\iota,a} \sim \mathcal{N}\left(-\frac{\sigma_\varepsilon^2}{2(1 + \eta_\varepsilon)}, \sigma_\varepsilon^2\right). \quad (28)$$

Here σ_ε is the standard deviation of the per-age innovation. The shock is initialized at the first working age a_{e_ι} from the stationary distribution of (27), so that, together with the negative innovation mean, the multiplicative shock has unit mean at every age, $\mathbb{E}[\exp(\varepsilon_{\iota,a})] = 1$.

Following Imai and Keane (2004), agents accumulate human capital through learning-by-

doing. Human capital at the age of education completion is initialized by education level,

$$h_{\ell, a_{e_\ell}} = h_{e_\ell}, \quad h_{e_b} = 1, \quad h_{e_h} = 1 + \Delta h_h, \quad h_{e_c} = 1 + \Delta h_h + \Delta h_c, \quad (29)$$

where the below-high-school level is normalized to $h_{e_b} = 1$ and the high-school and college premia $\Delta h_h, \Delta h_c > 0$. Define the normalized learning function

$$g_h(\ell) = \frac{(\ell + \mathfrak{d}_h)^{\alpha_h} - \mathfrak{d}_h^{\alpha_h}}{(1 + \mathfrak{d}_h)^{\alpha_h} - \mathfrak{d}_h^{\alpha_h}}, \quad \ell \in [0, 1], \quad \alpha_h \in (0, 1), \quad \mathfrak{d}_h > 0. \quad (30)$$

This normalization implies $g_h(0) = 0$ and $g_h(1) = 1$, while the smoothing parameter $\mathfrak{d}_h$ keeps the marginal benefit of working time finite at $\ell_{\ell, a} = 0$. Since $\alpha_h \in (0, 1)$, the marginal effect of working time on human-capital accumulation is decreasing in $\ell_{\ell, a}$.

For working ages $a \geq a_{e_\ell}$, human capital then evolves according to

$$h_{\ell, a+1} = [(1 - \delta_h) + \gamma_{h,0} [1 - \gamma_{h,1} (a - a_{e_\ell})] g_h(\ell_{\ell, a})] h_{\ell, a}, \quad (31)$$

where δ_h is the depreciation rate of human capital, $\gamma_{h,0}$ is the learning-by-doing scale, and $\gamma_{h,1}$ characterizes the decline in learning ability as agents age, measured by time since education completion.

The college full-time reference path $h_{max, a}$, used to normalize the human-capital density in (5), is the trajectory of a college-educated agent who works full-time ($\ell_{\ell, a} = 1$) at every age:

$$h_{max, a_{e_c}} = h_{e_c}, \quad h_{max, a+1} = [(1 - \delta_h) + \gamma_{h,0} (1 - \gamma_{h,1} (a - a_{e_c}))] h_{max, a}, \quad a \geq a_{e_c}, \quad (32)$$

with $h_{max, a} = h_{e_c}$ for the pre-working ages $a < a_{e_c}$.

Figure 15 summarizes the computation of human capital and labor supply.




Figure 15: Labor Supply and Human Capital Dynamics

Notes: Panel (a) shows the AR(1) idiosyncratic labor productivity shock. Panel (b) shows human-capital accumulation from current human capital and working time, with $h_{\ell, a}$ exogenously initialized at $a = a_{e_\ell}$. Panel (c) shows effective labor supply as a contemporaneous function of the shock, working time, and human capital.

2.5 Agents' Prediction System

Having specified agents' preferences, human capital accumulation, and labor supply, I now describe how agents form expectations and choose consumption, saving, and working time. As I show below, a rational-expectations benchmark is computationally prohibitive, so agents instead form predictions in a deterministic auxiliary economy, whose construction occupies the rest of this subsection. As a first step, I fix notation for the living population, their individual states, and the information set they share.

Let

$$\mathcal{L}_{j,t} := \{(i, j, t_0) : i \in I, t_0 \leq t \leq t_0 + d_{j,t_0} - 1\} \quad (33)$$

denote the set of agents alive in country j at time t . For any $\iota = (i, j, t_0) \in \mathcal{L}_{j,t}$, let $a_{\iota,t} := t - t_0 + 1$ be her age at time t . The individual state of agent ι at time t is

$$\mathbf{s}_{\iota,t} := (k_{\iota,a_{\iota,t}}, h_{\iota,a_{\iota,t}}, \varepsilon_{\iota,a_{\iota,t}}). \quad (34)$$

The entries in $\mathbf{s}_{\iota,t}$ are, respectively, last-period saving $k_{\iota,a_{\iota,t}}$ (the agent's start-of-period wealth, predetermined when she enters age a), human capital, and the idiosyncratic labor productivity shock.

The individual-state block collects these states across all living agents:

$$\mathbf{S}_{I,j,t} := \{\mathbf{s}_{\iota,t} : \iota \in \mathcal{L}_{j,t}\}. \quad (35)$$

Thus, $\mathbf{S}_{I,j,t}$ summarizes the cross-sectional distribution of predetermined saving, human capital, and productivity shocks over families and cohorts in country j .

The aggregate-state block contains variables common to all agents in the same country-year:

$$\begin{aligned} \mathbf{S}_{A,j,t} := & \left(\{(N_{j,t_0}, d_{j,t_0})\}_{t_0 \leq t}, \right. \\ & \left. \{(\underline{x}_{j,t_0}, \mathbf{e}_{h,j,t_0}, \mathbf{e}_{c,j,t_0}, \tau_{k,j,t_0}, \bar{\tau}_{l,j,t_0}, \tau_{m,j,t_0}, \zeta_{j,t_0})\}_{t_0 \leq t}, \right. \\ & \left. K_{s,j,t}, \lambda_{j,t}, U_{j,t} \right), \end{aligned} \quad (36)$$

The first line of $\mathbf{S}_{A,j,t}$ is the demographic history, including entrant cohort sizes and life expectancy. The second line is the government-policy history: the consumption floor, education policy, tax policy, and government efficiency. The final line gives current state-owned capital and the technology state, summarized by the TFP-to-ceiling ratio and the TFP shock. The

information set shared by all agents in country j at time t is

$$\mathbf{S}_{j,t} := (\mathbf{S}_{I,j,t}, \mathbf{S}_{A,j,t}). \quad (37)$$

Thus, $\mathbf{S}_{j,t}$ combines the cross-sectional distribution of individual states of all living agents with the aggregate states observed at time t .

At age a , the agent chooses consumption and working time. I write this action compactly as

$$\mathbf{a}_{\iota,a} := (c_{\iota,a}, \ell_{\iota,a}) \in \mathfrak{A}_{\iota,a}, \quad (38)$$

where $\mathfrak{A}_{\iota,a}$ is the feasible action set. The agent enters age a with start-of-period wealth $k_{\iota,a}$. Scaling this wealth by the contemporaneous after-tax gross return, adding current-age after-tax labor income $y_{\iota,a}$, and topping the total up to the consumption floor $\underline{x}_{j,t}$ gives her cash-on-hand $x_{\iota,a}$. What remains after consumption is her end-of-period saving $k_{\iota,a+1}$, which becomes next period's start-of-period wealth.

$$\begin{aligned} x_{\iota,a} &= \max\{(1 + r_{p,j,t} - \max\{0, \tau_{k,j,t} r_{p,j,t}\})k_{\iota,a} + y_{\iota,a}, \underline{x}_{j,t}\}, \\ k_{\iota,a+1} &= x_{\iota,a} - c_{\iota,a} \geq 0. \end{aligned} \quad (39)$$

The capital-income tax falls only on positive private-capital returns, as in (17). Feasibility requires $0 < c_{\iota,a} \leq x_{\iota,a}$ and $\ell_{\iota,a} \in [0, 1]$. For any $\mathbf{a} = (c, \ell)$, I also write $u_{\iota}(\mathbf{a})$ for $u_{\iota}(c, \ell)$.

With the state and action notation in place, a natural benchmark for modeling agent choices is rational expectations (RE). Under RE, an agent $\iota = (i, j, t_0)$ at age $a < d_{j,t_0}$ and calendar time $t = t_0 + a - 1$ tracks the transition law of the entire information set and solves

$$V_{\iota,a}(\mathbf{S}_{j,t}) = \max_{\mathbf{a} \in \mathfrak{A}_{\iota,a}} \left\{ u_{\iota}(\mathbf{a}) + \beta_i \int V_{\iota,a+1}(\mathbf{S}_{j,t+1}) p_{\mathbf{S}}(\mathbf{S}_{j,t}, \mathbf{a}, d\mathbf{S}_{j,t+1}) \right\}, \quad (40)$$

where $V_{\iota,a}(\cdot)$ is the value function at age a of agent ι , and $p_{\mathbf{S}}$ is the transition probability function for the information set $\mathbf{S}_{j,t}$. As Section 1 sets out, this Bellman equation is computationally prohibitive: the state is infinite-dimensional, because future prices depend on the full cross-sectional distribution of individual states, and the environment is non-stationary, so the economy follows a single calendar-time path rather than the recursive structure that standard solution methods exploit.

I therefore replace rational expectations equilibrium with a bounded-rationality concept

that I call the certainty-equivalent-anchored (CE-anchored) quasi-rational equilibrium.¹⁴ In each period, agents first construct a certainty-equivalent (CE) price path in a deterministic auxiliary economy. This CE path anchors a subjective stochastic price process, so every agent solves her own life-cycle problem with subjective price risk and idiosyncratic labor productivity risk rather than forecasting the full future distribution of all the state variables in the economy. At the start of time t , this procedure yields a subjective stochastic process for labor and private-capital prices over the current and future periods (the current price is not yet known when working time is chosen) together with predicted paths for the future tax-and-transfer system:

$$\{(\hat{w}_{j,t \rightarrow s}, \hat{r}_{p,j,t \rightarrow s})\}_{s \geq t}, \quad \{(\hat{\tau}_{l,j,t \rightarrow s}, \hat{\tau}_{k,j,t \rightarrow s}, \hat{\tau}_{m,j,t \rightarrow s}, \hat{x}_{j,t \rightarrow s})\}_{s \geq t+1}. \quad (41)$$

The notation of the form $\hat{x}_{j,t \rightarrow s}$ denotes agents' prediction at time t for the realization of a stochastic process $(x_{j,t})_{t \in \mathbb{N}}$ at a future time $s \geq t$. The colon convention for stacking these predictions over a horizon, $\hat{x}_{t \rightarrow t_1:t_2}$, follows [Section 2](#).

Agents form the predictions in (41) in three steps. First, agents forecast demography $\hat{N}_{j,t \rightarrow t+1:\infty}$ and the tax-and-transfer policies with simple extrapolation rules. Second, they construct a deterministic auxiliary economy in which TFP and individual shocks follow their certainty-equivalent (CE) paths; the wage and private-capital return that clear it over the forecast window $[t : t + \Delta]$ (the current period and the next Δ years, i.e. $\Delta + 1$ price pairs) are the CE price anchors:

$$(\bar{w}_{j,t \rightarrow t:t+\Delta}, \bar{r}_{p,j,t \rightarrow t:t+\Delta}) \quad (42)$$

Beyond the forecast horizon, these anchors stay fixed at their terminal-horizon values,

$$(\bar{w}_{j,t \rightarrow t'}, \bar{r}_{p,j,t \rightarrow t'}) = (\bar{w}_{j,t \rightarrow t+\Delta}, \bar{r}_{p,j,t \rightarrow t+\Delta}), \quad t' > t + \Delta.$$

Finally, they attach subjective price risk around these anchors, forming the stochastic price process that enters their individual decision problem. The six blocks below develop these steps in order.

Demography Prediction Let $\hat{N}_{j,t \rightarrow t'}$ denote the prediction formed in year t for the size of the age-1 entrant cohort in country j in year t' . For realized years, the prediction is initialized

¹⁴“Bounded rationality” has many interpretations. For example, [Krusell and Smith \(1998\)](#) uses it to describe agents who learn the laws of motion of the aggregate variables. Here, I use it simply to refer to constraints on an agent’s “computational capacity” in the sense of [Simon \(1955\)](#). These assumptions do not fully answer [Simon \(1996\)](#)’s call for empirical evidence on how bounded-rational actors form and use expectations, but they provide a tractable foundation for the analysis.

with the observed history:

$$\hat{N}_{j,t \rightarrow t'} = N_{j,t'}, \quad t' \leq t. \quad (43)$$

For future years, agents keep the current fertility index fixed. I define this index as current entrants relative to the lagged cohort base:

$$\varphi_{j,t}^N := \frac{N_{j,t}}{\sum_{a=1}^{a_{j,max,t}} N_{j,t-a}}, \quad (44)$$

The predicted entrant-cohort size then evolves recursively as

$$\hat{N}_{j,t \rightarrow t'} = \varphi_{j,t}^N \sum_{a=1}^{a_{j,max,t}} \hat{N}_{j,t \rightarrow t'-a}, \quad t' > t. \quad (45)$$

Thus, demographic prediction reduces to propagating the observed cohort history forward under a fixed fertility index. This frozen-rate rule is deliberately simple and follows the spirit of evidence that current fertility-rate extrapolations can be competitive with more elaborate fertility forecasts (Bohk-Ewald et al., 2018). Figure 16 summarizes the demographic forecasting rule.




Figure 16: Demography Prediction

Notes: The observed cohort history determines the fixed fertility index $\varphi_{j,t}^N$, which propagates entrant cohorts forward.

Tax and Transfer Policies Prediction Agents predict future tax policies by freezing the tax system at its current value. For any future year $t' > t$,

$$\begin{aligned} \hat{\tau}_{k,j,t \rightarrow t'} &= \tau_{k,j,t}, \\ \hat{\mathcal{T}}_{l,j,t \rightarrow t'}(\cdot) &= \mathcal{T}_{l,j,t}(\cdot), \\ \hat{\tau}_{m,j,t \rightarrow t'} &= \tau_{m,j,t}. \end{aligned} \quad (46)$$

Agents also keep fixed the current implied floor-to-wage ratio. Let

$$\hat{\Lambda}_{j,t} := \frac{\underline{x}_{j,t}}{w_{j,t}}. \quad (47)$$

Once the CE wage anchor $\bar{w}_{j,t \rightarrow t'}$ is formed, the predicted consumption floor is

$$\hat{\underline{x}}_{j,t \rightarrow t'} = \hat{\Lambda}_{j,t} \bar{w}_{j,t \rightarrow t'}, \quad t' \geq t. \quad (48)$$

Thus, agents allow the level of the consumption floor to move with the CE wage anchor, but they do not forecast changes in the fiscal rules or in the floor-to-wage ratio. Figure 17 summarizes the policy forecast.




Figure 17: Tax and Transfer Policy Prediction

Notes: Panel (a) freezes the tax-system tuple at its current value. Panel (b) keeps the floor-to-wage ratio fixed, so the predicted consumption floor moves with the CE wage anchor.

TFP Prediction Agents forecast TFP by removing future innovations from the TFP-shock process in (6). The forecast is initialized at the current shock,

$$\hat{U}_{j,t \rightarrow t} = U_{j,t},$$

and for any future year $t' > t$ agents set $\epsilon_{u,j,t'} = 0$, so that

$$\hat{U}_{j,t \rightarrow t'} = \eta_{j,u} \hat{U}_{j,t \rightarrow t'-1}. \quad (49)$$

Along the CE price anchors in (42), the CE agents' decision problem below determines the human-capital distribution, $\hat{\mathbf{h}}_{j,t \rightarrow t'}$. Predicted TFP is then obtained by applying the TFP system in (2)–(6) to this distribution and the predicted cohort sizes, with each agent's education setting her weight in the education-weighted density of (5). Thus, conditional on the CE price anchors, TFP prediction is deterministic. Figure 18 summarizes the TFP prediction block.




Figure 18: TFP Prediction

Notes: CE price anchors shape the predicted human-capital distribution, which also records each agent's education and hence education weight. Together with the predicted cohort sizes, it determines the education-weighted density, catch-up, and predicted TFP.

CE Agents' Optimization Problem The auxiliary economy is populated by CE agents who take the deterministic predicted environment as given. For $\iota = (i, j, t_0)$ alive at t , let $a = a_{\iota,t}$

be the age at which the prediction is formed. For any future age $a_1 \geq a$, let $t_1 = t_0 + a_1 - 1$. The predicted individual state

$$\hat{\mathbf{s}}_{\ell, a \rightarrow a_1} = (\hat{k}_{\ell, a \rightarrow a_1}, \hat{h}_{\ell, a \rightarrow a_1}, \hat{\varepsilon}_{\ell, a \rightarrow a_1}) \quad (50)$$

is the CE-path counterpart of the three components of (34), initialized at the current state, $\hat{\mathbf{s}}_{\ell, a \rightarrow a} = \mathbf{s}_{\ell, t}$. The agent values this state along a candidate future price path $\check{\mathbf{p}}_{j, t: \infty} = (\check{w}_{j, s}, \check{r}_{p, j, s})_{s \geq t}$, which is a sequence of wages and private-capital returns from age a onward. For $a_1 < d_{j, t_0}$, the CE value satisfies the deterministic Bellman equation

$$\bar{V}_{\ell, a \rightarrow a_1}(\hat{\mathbf{s}}_{\ell, a \rightarrow a_1}; \check{\mathbf{p}}_{j, t_1: \infty}) = \max_{\mathbf{a} \in \mathfrak{A}_{\ell, a_1}} \left\{ u_{\ell}(\mathbf{a}) + \beta_i \bar{V}_{\ell, a \rightarrow a_1+1}(\hat{\mathbf{s}}_{\ell, a \rightarrow a_1+1}; \check{\mathbf{p}}_{j, t_1+1: \infty}) \right\}, \quad (51)$$

where $\bar{V}_{\ell, a \rightarrow a_1}$ is the CE value at future age a_1 of the prediction formed at age a , evaluated at the predicted state $\hat{\mathbf{s}}_{\ell, a \rightarrow a_1}$ under the candidate price path $\check{\mathbf{p}}_{j, t_1: \infty}$. The price path enters (51) as a fixed argument, so the equation carries no integral; government policies and individual shocks are likewise fixed on the CE path. In particular, the labor-productivity innovation is set to zero, as for the TFP shock in (49):

$$\hat{\varepsilon}_{\ell, a \rightarrow a} = \varepsilon_{\ell, a}, \quad \hat{\varepsilon}_{\ell, a \rightarrow a_1+1} = \eta_{\varepsilon} \hat{\varepsilon}_{\ell, a \rightarrow a_1}. \quad (52)$$

The dynamic budget maps the state forward along the anchor path. With $\mathbf{a} = (c, \ell)$, the current-age after-tax labor income is

$$\hat{y}_{\ell, a \rightarrow a_1} = \check{w}_{j, t_1} q_{\ell, a_1} - \hat{\tau}_{\ell, j, t \rightarrow t_1}(\check{w}_{j, t_1} q_{\ell, a_1}), \quad (53)$$

where efficient labor supply q_{ℓ, a_1} follows (26) from ℓ , current human capital $\hat{h}_{\ell, a \rightarrow a_1}$, and the CE productivity $\hat{\varepsilon}_{\ell, a \rightarrow a_1}$. The cash-on-hand and next period's saving are

$$\begin{aligned} \hat{x}_{\ell, a \rightarrow a_1} &= \max\{(1 + \check{r}_{p, j, t_1} - \max\{0, \hat{\tau}_{k, j, t \rightarrow t_1} \check{r}_{p, j, t_1}\}) \hat{k}_{\ell, a \rightarrow a_1} + \hat{y}_{\ell, a \rightarrow a_1}, \underline{x}_{j, t \rightarrow t_1}\}, \\ \hat{k}_{\ell, a \rightarrow a_1+1} &= \hat{x}_{\ell, a \rightarrow a_1} - c \geq 0. \end{aligned} \quad (54)$$

Human capital updates through (31), and the next state is

$$\hat{\mathbf{s}}_{\ell, a \rightarrow a_1+1} = (\hat{k}_{\ell, a \rightarrow a_1+1}, \hat{h}_{\ell, a \rightarrow a_1+1}, \hat{\varepsilon}_{\ell, a \rightarrow a_1+1}).$$

At the terminal age $a_1 = d_{j,t_0}$ the continuation is the bequest motive,

$$\bar{V}_{\iota,a \rightarrow d_{j,t_0}}(\hat{\mathbf{s}}_{\iota,a \rightarrow d_{j,t_0}}) = \max_{\mathbf{a} \in \mathfrak{A}_{\iota,d_{j,t_0}}, \mathbf{g} \geq 0} \{u_{\iota}(\mathbf{a}) + u_{\iota}^{\text{beq}}(\mathbf{g})\} \quad \text{s.t.} \quad c + \mathbf{g} = \hat{x}_{\iota,a \rightarrow d_{j,t_0}}, \quad (55)$$

where $\hat{x}_{\iota,a \rightarrow d_{j,t_0}}$ is the terminal cash-on-hand of (54).

Solving the deterministic DP problem in (51) maps the current state and a price path into predicted paths for savings, human capital, and working time—written $\hat{k}$, $\hat{h}$, $\hat{\ell}$ —over the next Δ years. Figure 19 summarizes the CE individual problem.




Figure 19: CE Agent Optimization

Notes: The current individual state and CE price anchors determine the future path tuple used for aggregation. If the original agent dies within the forecast horizon, the tuple refers to the corresponding descendant.

Imagined Market Clearing For a candidate price path $\check{\mathbf{p}}_{j,t:t+\Delta}$ and the deterministic predictions above, the CE agents solve the CE problem (51)–(52) at $\check{\mathbf{p}}_{j,t:t+\Delta}$. A cohort entering after t (entry year $t_0 \in \{t+1, \dots, t+\Delta\}$) solves the same CE problem from age 1, initialized at its predicted age-1 state: the bequest left by its exiting ancestor (Section 2.1), predicted human capital, and the certainty-equivalent labor-productivity shock. The resulting individual paths imply aggregate private-capital and labor supplies for each year in the clearing window $t' \in \{t, \dots, t+\Delta\}$:

$$\begin{aligned} \hat{K}_{p,j,t \rightarrow t'} &= \sum_{a=1}^{a_{j,max,t'}} \hat{N}_{j,t \rightarrow t' - a + 1} \int \hat{k}_{(i,j,t' - a + 1),t \rightarrow a} \mathbf{P}_j(di), \\ \hat{L}_{j,t \rightarrow t'} &= \sum_{a=1}^{a_{j,max,t'}} \hat{N}_{j,t \rightarrow t' - a + 1} \int \hat{h}_{(i,j,t' - a + 1),t \rightarrow a} \hat{\ell}_{(i,j,t' - a + 1),t \rightarrow a} \exp\left(\hat{\varepsilon}_{(i,j,t' - a + 1),t \rightarrow a}\right) \mathbf{P}_j(di), \end{aligned} \quad (56)$$

as in (8) and (9). In the imagined world, the predicted tax-and-transfer policies, applied to the CE agents and prices, generate the capital (17), labor (18), and transfer (20) tax revenues. These, plus the gross return on existing state-owned capital and scaled by government fiscal efficiency, constitute predicted revenue, while predicted welfare (consumption-floor) transfers (13) and education spending (16) are the expenditure. The government budget (22) then leaves predicted state-owned capital $\hat{K}_{s,j,t \rightarrow t'}$ as the residual, so aggregate capital $\hat{K}_{j,t \rightarrow t'}$ follows from (7).

Given $\hat{A}_{j,t \rightarrow t'}$, $\hat{K}_{p,j,t \rightarrow t'}$, $\hat{K}_{s,j,t \rightarrow t'}$, and $\hat{L}_{j,t \rightarrow t'}$, auxiliary-economy output and CES shares are

$$\begin{aligned}\hat{Y}_{j,t \rightarrow t'} &= \hat{A}_{j,t \rightarrow t'} \left(\alpha_{Y,j} \hat{K}_{j,t \rightarrow t'}^{\rho_{Y,j}} + (1 - \alpha_{Y,j}) \hat{L}_{j,t \rightarrow t'}^{\rho_{Y,j}} \right)^{1/\rho_{Y,j}}, \\ \hat{s}_{p,j,t \rightarrow t'} &= \frac{\alpha_{K,j} \hat{K}_{p,j,t \rightarrow t'}^{\rho_{K,j}}}{\alpha_{K,j} \hat{K}_{p,j,t \rightarrow t'}^{\rho_{K,j}} + (1 - \alpha_{K,j}) \hat{K}_{s,j,t \rightarrow t'}^{\rho_{K,j}}}, \\ \hat{s}_{K,j,t \rightarrow t'} &= \frac{\alpha_{Y,j} \hat{K}_{j,t \rightarrow t'}^{\rho_{Y,j}}}{\alpha_{Y,j} \hat{K}_{j,t \rightarrow t'}^{\rho_{Y,j}} + (1 - \alpha_{Y,j}) \hat{L}_{j,t \rightarrow t'}^{\rho_{Y,j}}}, \quad \hat{s}_{L,j,t \rightarrow t'} = 1 - \hat{s}_{K,j,t \rightarrow t'}.\end{aligned}\tag{57}$$

These are the CE counterparts of (1) and (10). The firm-implied wage and private-capital return, following (11), are

$$\begin{aligned}w_{j,t \rightarrow t'}^f &= \frac{\hat{Y}_{j,t \rightarrow t'}}{\hat{L}_{j,t \rightarrow t'}} \hat{s}_{L,j,t \rightarrow t'}, \\ r_{p,j,t \rightarrow t'}^f &= \frac{\hat{Y}_{j,t \rightarrow t'}}{\hat{K}_{p,j,t \rightarrow t'}} \hat{s}_{K,j,t \rightarrow t'} \hat{s}_{p,j,t \rightarrow t'} - \delta,\end{aligned}\tag{58}$$

These firm-implied prices depend on the candidate $\check{\mathbf{p}}_{j,t:t+\Delta}$ through the CE agents' decisions. The CE anchor $\bar{\mathbf{p}}_{j,t:t+\Delta} := (\bar{w}_{j,t:t+\Delta}, \bar{r}_{p,j,t:t+\Delta})$ is the candidate at which the auxiliary economy comes closest to market clearing. Collecting a candidate over the clearing window as $\check{\mathbf{p}}_{j,t:t+\Delta}$, the L_2 discrepancy between the candidate and the firm-implied prices it generates is

$$\mathcal{J}_{j,t}(\check{\mathbf{p}}_{j,t:t+\Delta}) := \sum_{t'=t}^{t+\Delta} \left[\left(\check{w}_{j,t'} - w_{j,t \rightarrow t'}^f \right)^2 + \left(\check{r}_{p,j,t'} - r_{p,j,t \rightarrow t'}^f \right)^2 \right],\tag{59}$$

and the anchor is its minimizer,

$$\bar{\mathbf{p}}_{j,t:t+\Delta} = \underset{\check{\mathbf{p}}_{j,t:t+\Delta}}{\operatorname{argmin}} \mathcal{J}_{j,t}(\check{\mathbf{p}}_{j,t:t+\Delta}).\tag{60}$$

Figure 20 summarizes the L_2 loss computation.




Figure 20: Imagined Market Clearing

Notes: The information set and a candidate price path determine aggregate private capital, government capital, and labor over the forecast horizon. Firms imply demand-side prices from these aggregates, and the CE anchor minimizes the L_2 loss between the candidate and firm-implied prices.

Predicted Prices Process The CE price anchors are deterministic. For any year $t' \geq t$, the stochastic price process used in agents' individual decision problem scales those anchors by a

multiplicative log-normal price-risk factor:

$$\hat{w}_{j,t \rightarrow t'} = \bar{w}_{j,t \rightarrow t'} \exp(\xi_{w,j,t \rightarrow t'}), \quad \hat{r}_{p,j,t \rightarrow t'} = \bar{r}_{p,j,t \rightarrow t'} \exp(\xi_{r,j,t \rightarrow t'}). \quad (61)$$

The log price-risk terms, collected in $\boldsymbol{\xi}_{j,t \rightarrow t'} = (\xi_{w,j,t \rightarrow t'}, \xi_{r,j,t \rightarrow t'})$, each follow an AR(1) process:

$$\xi_{w,j,t \rightarrow t'} = \eta_{w,j} \xi_{w,j,t \rightarrow t'-1} + \epsilon_{w,j,t'}, \quad \xi_{r,j,t \rightarrow t'} = \eta_{r,j} \xi_{r,j,t \rightarrow t'-1} + \epsilon_{r,j,t'}, \quad (62)$$

with initial conditions $\xi_{w,j,t \rightarrow t} = \xi_{r,j,t \rightarrow t} = 0$, i.e., the current-period predicted prices coincide exactly with the anchors, with price risk entering only for future periods $t' > t$. The innovations to labor and private-capital prices are jointly Gaussian and may be correlated:

$$\epsilon_{w,j,t'} \sim \mathcal{N}(0, \sigma_{w,j}^2), \quad \epsilon_{r,j,t'} \sim \mathcal{N}(0, \sigma_{r,j}^2), \quad \text{Corr}(\epsilon_{w,j,t'}, \epsilon_{r,j,t'}) = \rho_{wr,j}. \quad (63)$$

Here $\sigma_{w,j}$ and $\sigma_{r,j}$ are the standard deviations of the log price-risk innovations, and $\rho_{wr,j}$ is the correlation between the wage and private-return innovations. Figure 21 summarizes the stochastic price process.




Figure 21: Predicted Stochastic Price Process

Notes: Panel (a) shows the bivariate AR(1) subjective price-risk process with correlated wage and private-return innovations. Panel (b) scales the deterministic CE anchor by the realized log-normal price-risk factor to form the predicted price distribution at t' .

2.6 Agents' Dynamic Optimization

The previous subsection equips each agent with a predicted environment: a fixed path of future government policies and a subjective price process. Even in this simplified setting, any agent ι of age a at time t still faces a stochastic DP problem with five state variables (saving, human capital, the labor-productivity shock, and the price-risk terms for the wage and the interest rate):

$$\hat{V}_{\iota,a}(\mathbf{s}_{\iota,a}, \boldsymbol{\xi}_{j,t \rightarrow t}) = \max_{\mathbf{a} \in \mathfrak{A}_{\iota,a}} \left\{ u_{\iota}(\mathbf{a}) + \beta_i \iint \hat{V}_{\iota,a+1}(\mathbf{s}_{\iota,a+1}, \boldsymbol{\xi}_{j,t \rightarrow t+1}) \right. \\ \left. \hat{p}_{\iota,t}(\mathbf{s}_{\iota,a}, \mathbf{a}, d\mathbf{s}_{\iota,a+1} \mid \boldsymbol{\xi}_{j,t \rightarrow t+1}) dF_{\boldsymbol{\xi}}(\boldsymbol{\xi}_{j,t \rightarrow t+1}) \right\}. \quad (64)$$

The inner integral, against the individual-state kernel $\hat{p}_{l,t}$, averages over the idiosyncratic productivity innovation; the outer integral, against F_{ξ} , averages over the price-risk innovation of (62).

Solving (64) exactly is another obstacle for agents' decision, beyond the infinite-dimensional state and the non-stationarity that ruled out the rational-expectations benchmark (40): the five-dimensional non-convex DP problem is computationally demanding, and it needs to be solved for every agent type, age, and year at every parameter draw the estimation visits. I therefore assume the agent makes two simplifications.

First, the agent values the future not by the unknown stochastic continuation $\hat{V}$ but by the certainty-equivalent continuation $\bar{V}$ of (51). Conditional on a realized price-risk vector $\xi_{j,t \rightarrow t+1}$, the future price path is deterministic—the anchors tilted by $\xi_{j,t \rightarrow t+1}$ and decaying back through (62) with the innovations switched off—so the agent evaluates the certainty-equivalent value $\bar{V}_{l,a+1 \rightarrow a+1}(\mathbf{s}_{l,a+1}; \mathbf{p}(\xi_{j,t \rightarrow t+1}))$ on that tilted anchor path $\mathbf{p}(\xi_{j,t \rightarrow t+1})$ —the year- t anchors of (42) tilted by the draw, reused as the continuation's belief rather than re-cleared at $t + 1$:

$$\begin{aligned} \mathbf{p}(\xi_{j,t \rightarrow t+1}) &= \{(w_{j,t \rightarrow t'}, r_{p,j,t \rightarrow t'})\}_{t' \geq t+1}, \\ w_{j,t \rightarrow t'} &= \bar{w}_{j,t \rightarrow t'} \exp\left(\eta_{w,j}^{t'-t-1} \xi_{w,j,t \rightarrow t+1}\right), \\ r_{p,j,t \rightarrow t'} &= \bar{r}_{p,j,t \rightarrow t'} \exp\left(\eta_{r,j}^{t'-t-1} \xi_{r,j,t \rightarrow t+1}\right). \end{aligned} \tag{65}$$

Second, the agent takes a single improvement step instead of iterating to a fixed point. The implemented decision rule is the one-step rollout

$$\begin{aligned} \hat{\pi}_{l,a}(\mathbf{s}_{l,a}) \in \arg \max_{\mathbf{a} \in \mathfrak{A}_{l,a}} \left\{ u_l(\mathbf{a}) + \beta_i \iint \bar{V}_{l,a+1 \rightarrow a+1}(\mathbf{s}_{l,a+1}; \mathbf{p}(\xi_{j,t \rightarrow t+1})) \right. \\ \left. \hat{p}_{l,t}(\mathbf{s}_{l,a}, \mathbf{a}, d\mathbf{s}_{l,a+1} \mid \xi_{j,t \rightarrow t+1}) dF_{\xi}(\xi_{j,t \rightarrow t+1}) \right\}, \end{aligned} \tag{66}$$

evaluated at the current state and re-formed each period as new states realize. The inner kernel averages the certainty-equivalent continuation over the next individual state—the year- $(t + 1)$ labor productivity draw, and the outer integral averages over the next-period price risk $\xi_{j,t \rightarrow t+1}$, which equals the one-step innovation of (63) since the year- t price risk is zero, $\xi_{j,t \rightarrow t} = \mathbf{0}$ (Section 2.5); this draw also tilts the continuation's price path $\mathbf{p}(\xi')$. Relative to (64), the rule substitutes the certainty-equivalent continuation $\bar{V}$ for the stochastic $\hat{V}$ and takes a single Newton step improvement on (64) over the certainty-equivalent policy (Bertsekas, 2025). In computing $\bar{V}$, I assume agents choose a labor-supply path that is a logit transformation of a quadratic in age, which simplifies their computation in Section 3.3.

The one-step rollout is both economically attractive and computationally feasible. It is attractive because it still integrates over the next-period productivity and price-risk draws—preserving much of the precautionary-saving motive—and because it is a one-step policy improvement on the certainty-equivalent policy. It is feasible because it needs only the certainty-equivalent continuation $\bar{V}$ and the year- $(t + 1)$ integral, so each choice reduces to a static maximization rather than a backward induction over the five-dimensional stochastic state of (64).

Figure 22 visualizes the one-step rollout decision rule (66). Starting from current individual state $\mathbf{s}_{\ell,a}$, the agent first picks a consumption and working-time action $\mathbf{a}$, and the fan of next-period states then branches once per integral, first over the price risk ξ' and then over the labor-productivity draw $\epsilon_{\varepsilon,\ell,a+1}$. The right panel values each resulting state by the certainty-equivalent continuation along the anchor path its own price draw tilts, and the agent keeps the action whose payoff plus β_i -discounted average continuation is largest.




Figure 22: The One-Step Rollout Decision Rule

Notes: The figure schematizes (66), and its layout is inspired by Bertsekas (2025, Figure 1.6.12). The fan branches once per integral: the outer branch is the price risk $\xi' \equiv \xi_{j,t \rightarrow t+1}$, and the inner branch is the labor-productivity innovation $\epsilon_{\varepsilon,\ell,a+1}$. Each solid dot is a next individual state $\mathbf{s}_{\ell,a+1}$ drawn from the kernel $\hat{p}_{\ell,t}$, and the open circles are the ages the continuation runs through.

Figure 23 summarizes how the bounded-rationality construction makes the agent's problem tractable.




Figure 23: Agent Choice under Rational Expectations and Bounded Rationality

Notes: Carrying the full information set $\mathbf{S}_{j,t}$ in panel (a) requires the agents to track the whole economy. The predicted individual transition kernel in panel (b) makes the agent's problem manageable. The implemented year- t action evaluates the rollout decision rule (66) at the observed individual state, with the price-risk vector set to zero.

Under rational expectations, the agent maps the full information set directly into an action, so agents must track the transition law of the entire economy. Under bounded rationality, agents first form a predicted environment—price processes, tax-transfer policy paths, and individual transition kernels—and then choose by a one-step rollout over the certainty-equivalent value, rather than from the full aggregate state. The next subsection embeds these decision rules in the equilibrium.

2.7 Certainty-Equivalent-Anchored Quasi-Rational Equilibrium

Based on the agents' prediction system in [Section 2.5](#), this section defines the equilibrium concept, which I term **certainty-equivalent-anchored (CE-anchored) quasi-rational equilibrium**.

This approach creates a *two-layer reality*. First, agents construct an imagined deterministic world in which the price anchor $\bar{\mathbf{p}}_{j,t \rightarrow t:t+\Delta} := (\bar{w}_{j,t \rightarrow t:t+\Delta}, \bar{r}_{p,j,t \rightarrow t:t+\Delta})$ is found. Second, there is an *actual world* where unforeseen shocks occur, and outcomes diverge from the forecasts. Each period, agents observe the new state of the actual world and create a new forecast and re-solve their life-cycle problem.

Fix a country j and the structural parameters of the economy. These parameters include the forecast horizon Δ . A CE-anchored quasi-rational equilibrium is a collection of:

- (i) actual information sets $\{\mathbf{S}_{j,t}\}_{t \geq t_{s,j}}$, as defined in [\(37\)](#);
- (ii) actual prices $\{w_{j,t}, r_{p,j,t}, r_{s,j,t}\}_{t \geq t_{s,j}}$, determined by [\(1\)](#), [\(10\)](#), and [\(11\)](#);
- (iii) for each forecast year t , CE price anchors $\bar{\mathbf{p}}_{j,t \rightarrow t:t+\Delta} = \{(\bar{w}_{j,t \rightarrow s}, \bar{r}_{p,j,t \rightarrow s})\}_{s=t}^{t+\Delta}$, as in [\(42\)](#);
- (iv) for each forecast year t , subjective price-risk processes $\{\xi_{j,t \rightarrow s}\}_{s \geq t}$ and predicted price processes $\{(\hat{w}_{j,t \rightarrow s}, \hat{r}_{p,j,t \rightarrow s})\}_{s \geq t}$, as defined in [\(62\)](#) and [\(61\)](#).

These objects satisfy the following conditions:

1. **Initial condition.** The initial information set $\mathbf{S}_{j,t_{s,j}}$ is given and contains the individual and aggregate states required by [\(37\)](#).
2. **CE auxiliary-economy clearing.** In each year t , agents use $\mathbf{S}_{j,t}$ to form the deterministic demographic, policy, and TFP forecasts in [Section 2.5](#). The CE price anchors $\bar{\mathbf{p}}_{j,t \rightarrow t:t+\Delta}$ minimize the imagined auxiliary economy's clearing loss over the window $[t : t + \Delta]$, [\(59\)](#) and [\(60\)](#).

3. **Subjective price prediction.** Around these CE anchors, agents form subjective price risk and predicted stochastic prices according to (61) and (62).
4. **Individual optimization.** Given the predicted stochastic environment, each living agent forms an estimated transition kernel $\hat{p}_{\iota,t}$ for the one-step transition of her individual state (34), averaging over the year- $(t+1)$ labor-productivity draw. The agent implements the one-step rollout decision rule (66), which values the future by the certainty-equivalent continuation $\bar{V}$ and takes a single improvement step.
5. **Actual allocation and price determination.** The actual economy then evolves under the realized shock collection

$$\Xi_{j,t+1} := \left(\epsilon_{u,j,t+1}, \{ \epsilon_{\varepsilon,\iota,a_{\iota,t+1}} \}_{\iota \in \mathcal{L}_{j,t+1}} \right), \quad (67)$$

which stacks the aggregate TFP innovation $\epsilon_{u,j,t+1}$ of (6) and the idiosyncratic labor-productivity innovations $\epsilon_{\varepsilon,\iota,a_{\iota,t+1}}$ of (28) for every working-age agent alive at $t+1$. Individual states update through the realized budget and state law (39), human capital evolves by (31), the government budget updates by (22), aggregate private capital and labor are formed by (8)–(9) and combined into effective aggregate capital by (7), TFP follows (2)–(6), and realized output and factor prices are determined by (1), (10), and (11).

6. **Recursive updating.** In year $t+1$, agents observe the new information set $\mathbf{S}_{j,t+1}$ and repeat the same forecast–decision–realization cycle.

Figure 24 summarizes the recursive evolution of the actual economy. Starting from initial economy $\mathbf{S}_{j,t_{s,j}}$, agents form the CE price anchor $\bar{\mathbf{p}}_{j,t_{s,j} \rightarrow t_{s,j}+1:t_{s,j}+\Delta}$ and make decisions. Together with the shock collection $\Xi_{j,t_{s,j}+1}$ in (67), these decisions determine the economy at next period $\mathbf{S}_{j,t_{s,j}+1}$, and the cycle drives the economy forward — a process less of “Veni, vidi, vici” and more of “I predict, I act, I update”.




Figure 24: Evolution of the Economy

Notes: $\mathbf{S}_{j,t_{s,j}}$ is the initial actual information set, $\bar{\mathbf{p}}_{j,t_{s,j} \rightarrow t_{s,j}+1:t_{s,j}+\Delta}$ is the corresponding CE price anchor, and $\Xi_{j,t_{s,j}+1}$ collects all individual and aggregate shocks realized between $t_{s,j}$ and $t_{s,j}+1$.

3 Computation and Estimation

This section describes how I compute and estimate the model.¹⁵ The guiding idea is amortization: wherever the simulation would otherwise call an expensive solver repeatedly, I train a network to return that solver’s answer in a single forward pass. After constructing the data (Section 3.1) and the priors (Section 3.2), I amortize three solvers in turn. The agents’ dynamic programming problem becomes a forward pass through networks trained on solved decision rules (Sections 3.3 and 3.4). The anchor price path becomes a forward pass through a network trained without labels by descending the market-clearing loss itself, where needed, corrected by one learned Newton step in simulation (Section 3.5). The posterior becomes a forward pass through an inference network trained on simulated data, so inference without a tractable likelihood reduces to evaluating that network at the observed series (Section 3.6). Section 3.7 reports the resulting estimates. Throughout, the population is represented by 100 Monte Carlo families in both countries.

3.1 Data

I take the model to data on China (1982–2019) and Korea (1962–2019), assembling annual series on demographics, educational attainment, aggregate output and capital, a world technology frontier proxied by US total factor productivity, income and wealth inequality, and labor-income tax schedules. These draw on the World Bank, national population censuses, OUR WORLD IN DATA, CEIC, the National Bureau of Statistics, Barro and Lee (2013), the PENN WORLD TABLE (Feenstra et al., 2015), the WORLD INEQUALITY DATABASE (Chancel et al., 2022; Piketty et al., 2019; Hong et al., 2025), and the two countries’ income-tax policies. I normalize each country’s base-year population to one, all monetary variables are in constant 2021 USD. Appendix A describes the data sources and normalization in detail.

3.2 Prior Distributions of Parameters

As emphasized in Section 2, the priors are an integral part of the Bayes net. There are 68 estimated parameters in the model. Table 1 lists each one with its prior, set beside the posterior reported in Section 3.7 so that they can be read together.

The detailed rationale for each prior is given in Appendix B.

¹⁵All computations in this section run on a single laptop with 24 CPU threads and one NVIDIA GeForce RTX 5090 Laptop GPU. The code is written in Julia. The neural networks are implemented in `Lux.jl`; every training run is compiled to the GPU with `Reactant.jl` (XLA) and differentiated with `Enzyme.jl`, while the training labels are generated on CPU threads.

3.3 CE Agents' Decision Rules

The CE agents' decision rules are used repeatedly in the equilibrium: they shape both the agents' predictions of the anchor prices (Section 2.5–Section 2.7) and the agents' own decisions through the one-step rollout (Section 2.6). Re-solving them for every agent at every state would be prohibitively costly, so I approximate them with neural networks trained in two stages of *behavior cloning*. First, the neural networks of CE agents' saving, working-time, and value functions are fit to the optimal decisions and values that the Policy Tree Descendant (PTD) algorithm (Liu and Iskhakov, 2025), with a Nelder–Mead search, computes offline. Second, the rollout policy network is fit to the exact maximizers of the one-step rollout value, computed offline by a Nelder–Mead search with the pre-trained CE value function as its continuation (Section 3.4). See Amos (2023) for a general introduction to regression-based amortization. Figure 25 summarizes the pipeline.




Figure 25: Learning the CE Agents' Decision Rules

Solving CE Agents' Decision Rules The pipeline begins by generating training labels: for a large batch of CE agents, I solve the dynamic optimization problem (51) in three steps. I first parameterize the working-time path as a logit transformation of a quadratic in age, then compute the optimal lifetime discounted consumption-saving value for any given working-time path, and finally search over the three working-time parameters by Nelder-Mead algorithm to maximize this value.

The second step is the difficult one, because the consumption-saving problem is nonconvex: the consumption floor in (54) allows the agent save nothing. I solve it with the PTD algorithm. Conditional on a working-time path, the labor disutility in (24) is an additive constant and after-tax labor income is a known sequence, so (51) collapses to a deterministic consumption-saving problem. PTD represents its consumption policy function as a collection of affine segments $c_{\iota,a}^k(x) = b_0^k + b_1^k x$ over wealth intervals indexed by k . Its key property is that an affine policy

segment at age $a + 1$ has an affine descendant at age a under a linear Euler equation, which is ensured by the CRRA utility function. Thus, the policy function at every age can be traced by a recursive tree whose nodes are affine policy segments.

I illustrate the algorithm with an agent i at three periods from the exit age $d = 50$. The example uses a constant gross return $1 + r$ with $r = 0.1$, constant net labor income $y = 2.4$ per period, consumption floor $\underline{x} = 4.5$, discount factor $\beta_i = 0.98$, log utility ($\gamma_i = 1$), bequest weight $\psi_i = 0.1$, and no taxes. At the terminal age, the interior first-order condition of the bequest problem (55) gives a single affine segment $c_{i,50}(x) = x/(1 + \psi_i^{1/\gamma_i})$.

Stepping back to age 49, the Euler equation maps the terminal segment into one interior descendant $c_{i,49}^2(x)$, which competes with the save-nothing corner $c_{i,49}^1(x) = x$, under which next period's wealth falls to the consumption floor; Figure 26 (a) shows the two candidates. Each candidate implies a value once optimal behavior follows at age 50, and the two value segments cross at a single wealth level, marked by the black dot in panel (b). Keeping each candidate where it dominates gives a policy that is piecewise affine with one discontinuity, plotted in panel (c): below it, the agent consumes all wealth; above it, she saves along the interior segment.




Figure 26: Policy Tree Descendant: Backward Step from Age 50 to 49

Notes: (a) Candidate policy segments. (b) Value implied by each segment; the black dot marks the kink. (c) The refined policy function with one discontinuous jump.

The refined age-49 policy has two segments, each of which spawns its own interior descendant; with the save-nothing corner, three candidates now compete at age 48, as Figure 27 (a) shows. Fusing them pairwise through their value segments produces two kinks, the black dots in panel (b). The refined policy in panel (c) is piecewise affine with two discontinuities.




Figure 27: Policy Tree Descendant: Backward Step from Age 49 to 48

Notes: (a) Candidate policy segments. (b) Value implied by each segment; the black dots mark the two kinks. (c) The policy function with two discontinuous jumps.

Figure 28 collects the segments of the three example ages into the policy tree. Backward induction then continues down to the agent's current age, and a forward pass on the completed policy tree gives the exact consumption-saving plan and value for the given working-time path.




Figure 28: Policy Tree of the Example's Three Ages

Notes: Nodes are affine policy segments $c_{l,a}^k$; each orange dashed frame groups one age's segments. A solid arrow is the Euler interior descendant of a segment; the dashed nodes are the save-nothing corner.

PTD algorithm is grid-free: it carries the policy in closed form as affine segments and locates the discontinuities exactly, so cost scales with the number of segments rather than with the size of a wealth grid. Thus, it is fast and accurate. For an agent with exit age $d = 50$, one solve takes 0.01 to 1 millisecond (mean: 0.1 millisecond) on a single CPU thread. With the Nelder–Mead search over the three working-time parameters, one complete CE agent problem takes 2 to 300 milliseconds (mean: 40 milliseconds).

Learning CE Agents' Decision Rules PTD solves a single agent's CE problem exactly, but estimation computes the CE auxiliary-economy clearing millions of times, and each clearing

evaluates CE decision rules across the whole agent cross-section, billions of agents’ CE problems in all. I therefore amortize the solver into the three surrogate neural networks of Figure 25: the value net approximates the CE value $\bar{V}$ of (51), the saving net predicts the saving-rate path over the clearing window $[t : t + \Delta]$, and the labor net predicts the three working-time parameters. Each is a residual multi-layer perceptron (MLP) over the same feature vector encoding the agent’s state and CE environment: a dense layer mapping the feature vector to the hidden width, a stack of residual blocks whose identity skip connections keep gradients well-conditioned at depth, and a task-specific head as shown in Figure 29.




Figure 29: Residual MLP Architecture of the CE Agent Nets

Notes: Rectangles are dense layers; rounded boxes are residual blocks; each red arc (+) is the identity skip connection adding a block’s input to its output. The value and labor nets stack the three residual blocks drawn; the saving net inserts a fourth block and a compression layer before its head. The head returns the scalar value, the saving-rate path over the clearing window, or the three working-time parameters.

The labeled data set consists of 1,000,000 sampled CE problems labeled by the PTD solver with the Nelder–Mead search, a run of about two hours on 24 CPU threads, split by agent into 800,000 training, 100,000 validation, and 100,000 test problems. Each problem is one draw from the model’s own generative structure: a country and a calendar year fix the observed demography, education shares, and lifespans; the parameters come from the priors; the anchor price path is rebuilt from that year’s observed output and capital through the firm’s first-order conditions, perturbed to widen the coverage beyond the observed series; and the agent’s state is drawn to cover the reachable space, with initial wealth an income-scaled mixture of a zero atom, a constrained-poor uniform, a log-normal body, and a Pareto tail.

The labor and saving nets are trained on realized life-cycle path targets, while the value net regresses the optimal value compressed by the inverse hyperbolic sine (asinh). All three nets are trained with AdamW and early stopping on the validation loss; the full training run takes about eighteen minutes on the GPU. Figure 30 shows the per-epoch training and validation losses.

The trained nets are accurate on the held-out test problems. The saving net’s rate MAE is 0.013 on the test set. The labor net’s working-time error, scored on the test problems, is 0.011 per year. The predicted plan beats a random plan for 99.1% of 2,000 fresh agents. On the test problems, the value net attains an R^2 of 0.999 on the asinh scale and a median absolute error of

0.06 in value units. With the three nets in hand, evaluating a CE agent’s decision rule collapses from a full dynamic-programming solve to a single forward pass on the residual MLPs.




Figure 30: Training History of the CE Agent Nets

Notes: Per-epoch training (blue) and validation (orange) MSE of the labor (left), saving (center), and value (right) nets. The dashed line marks the best-validation epoch, the early-stopping checkpoint that is saved.

Appendix C details the sampler, input features, architectures, loss functions, and optimization.

3.4 Rollout Agents’ Decision Rules

Building on the CE agent nets, the rollout policy net amortizes the one-step rollout decision rule (66), which guides an agent’s current saving–working choices once she has formed the prediction of the future price process. Figure 31 draws the architecture.




Figure 31: MLP Architecture of the Rollout Policy Net

Notes: Rectangles are dense layers; the apricot box is the input feature vector, shared with the CE agent nets. The two heads are the two sigmoid units of the output layer, one per current-age control; the intervals are the bands the outputs are clamped to, the saving rate out of cash-on-hand and the working time ℓ .

Like the CE agent nets, the rollout policy net is trained by behavior cloning, on labels that solve the rollout problem exactly. For each of 100,000 sampled states, the two-control maximization of (66) is solved offline: a coarse grid over the control bands supplies global starts, and a derivative-free Nelder–Mead simplex refines them, with candidates projected onto

the bands inside the objective. The objective it maximizes, the braced term of (66) drawn in Figure 22, is evaluated by a 27-node Gauss–Hermite quadrature over the one-step innovations, with the pre-trained CE value net supplying the continuation value $\bar{V}_{t,a+1 \rightarrow a+1}$ at every node.¹⁶

The training states are drawn by the CE problem sampler of Section 3.3; labeling them takes about 75 minutes on 24 CPU threads, in chunks of 5,000 saved as resumable part files. Fitting the policy net is then a small supervised regression, mean squared error on the two controls, which early-stops in a few minutes (Figure 32).




Figure 32: Training History of the Rollout Policy Net

Notes: Per-epoch training (blue) and validation (orange) MSE of the supervised rollout policy net on the exact argmax labels: total (left), saving-rate head (center), and working-time head (right). The dashed line marks the best-validation epoch, the early-stopping checkpoint that is saved.

To measure the net’s out-of-sample performance, I re-solve the two-control maximization exactly on 5,000 fresh test states, by a dense grid search with local refinement. The net’s median relative optimality gap is 0.05%, and 94.5% of states fall within 1% of their optimum. Its controls sit further from the argmax than those gaps suggest, with MAE of 0.09 (medians 0.05) for the saving rate and 0.05 for the working time (median 0.02). For many agents the objective is flat in the saving rate, so a moderate distance from the argmax costs almost nothing in value. Appendix D details the rollout policy net architecture, the label solver, and the construction of the test metric.

3.5 Anchor Prices

The agents’ prediction system rests on the CE anchor path (42): the wage and private capital return path that brings the imagined auxiliary economy closest to market clearing, the minimizer (60) of the clearing loss (59). The simulation re-forms this anchor for every parameter draw

¹⁶Three Gauss–Hermite nodes per shock ($3^3 = 27$) integrate the smooth CE continuation values essentially exactly: against a 15^3 -node reference, the objective at the policy’s controls has a median relative error below 10^{-7} .

and every simulated year, so I amortize this solver as well. The anchor net maps an economy’s features, summarized by a small convolutional encoder, to the anchor path over the clearing window $[t : t + \Delta]$.

Unlike the agent and rollout nets, the anchor net is trained without labels: mini-batch gradient descent on the sampled-economy mean of the clearing loss (59) itself. Each training economy is itself a prior draw. Figure 33 draws the training graph, the full imagined-clearing forward pass of Figure 20: the frozen CE saving and labor nets of Section 3.3 choose saving and working time at the candidate prices, the agents’ states advance year by year over the clearing window, firms imply prices from the aggregated capital and labor supplies, and the candidate is scored by its L_2 discrepancy, summed over the window. The gradient is again the exact reverse-mode derivative of this loss, flowing back through the whole sweep, the frozen agent nets, and the convolutional encoder.




Figure 33: Objective-Based Training of the Anchor Net

Notes: Salmon boxes are the fixed, differentiable operations of the clearing loss (59); blue boxes are neural nets; apricot boxes are the objective’s inputs and the green box is the objective itself. The anchor net embeds the agent cross-section states through a convolutional encoder, joins it with the flat economy features, and decodes its head around the data-anchored warm start into the candidate anchor path.

The architecture of the anchor price net is drawn in Figure 34. The *cross-section grid* is the raw $14 \times 70 \times 100$ tensor of all individual agent states in the economy. The *flat features* are the 176 economy-level aggregate states. The CNN encoder turns the grid into a fixed-length summary. A cell MLP embeds every agent’s states into 16 channels with shared weights. I keep the sequence over ages, and because the age axis is ordered, two width-3 convolutions run along it: each output age reads its neighbors, letting a cohort’s summary reflect the cohorts entering around it. A masked mean-and-max pool over the alive ages then collapses the sequence into

a 64-dimensional embedding. That embedding joins the flat features at a residual block trunk of the same form as the CE agent nets (Figure 29), here of width 512. Its head returns 22 numbers, a wage and a return deviation for each of the eleven window offsets, and these are decoded about the data-anchored warm start rather than read as levels. The warm start is the firm-side price vector (11) evaluated at the observed output and capital of that country-year. In simulation, it is the previous year’s realized prices instead, held flat over the window.




Figure 34: Architecture of the Anchor Net

Notes: Apricot boxes are the net’s inputs. Blue boxes are learned layers. Salmon boxes are the pooling layers and the decoder. The last line of each block is the shape it outputs; parenthesized numbers are widths.

The anchor net is trained on 100,000 sampled economy-years. Figure 35 shows the training history on 255 held-out samples.




Figure 35: Training History of the Anchor Net

Notes: Left: per-step training clearing loss (blue, rescaled) and the held-out median (orange), on a $\log_{10}$ scale. Center and right: held-out median residuals, the largest relative wage gap and the largest absolute return gap between the candidate anchor and the firm-implied prices over the Δ -period clearing window. The dashed line marks the best-validation step.

The median clearing loss falls from 40.74 at the data-anchored warm start to 3.98 for the validation set, an order-of-magnitude improvement; in price units, the median largest relative wage gap over the clearing window falls from 50% to 19%, and the largest absolute return gap

from 7.7 to 2.1 percentage points.

Figure 36 collects the other parts that go into the anchor price amortizer: the objective-trained net is sharpened on an on-policy harvest, deployed in simulation, and backstopped there by a learned corrector.




Figure 36: Construction of the Anchor Price Amortizer

Notes: Blue boxes are learned objects, the salmon box is the deterministic simulation, and apricot boxes are the data each learned stage is trained on. Solid arrows run forward through the workflow; the dashed arrow is the gate, which sends only the years whose clearing residual exceeds tolerance to the corrector.

As shown in Figure 36, the anchor price net alone is not enough in simulation: on a held-out split of 4,463 economy-years harvested from simulator runs, its median relative wage gap is about 20%. Supervised fine-tuning on the other 40,113, each labeled by the exact anchor solve (two damped Gauss–Newton steps on the clearing loss), cuts the gap on those same held-out economy-years to 4.6%. This sharpened net is the one deployed.

It goes no further because the Gauss–Newton step that would close the gap needs the Jacobian of the clearing residual with respect to the heads, which no summary of the economy’s features supplies. Thus, I make the stage *semi-amortized* (Amos, 2023): a second small network, the *corrector*, distills the exact damped Gauss–Newton update on the market clearing loss (59), mapping the amortized prices, the measured clearing residual, and the flat features to a price vector correction, so that the probe survey gives way to a single clearing evaluation. It is trained on another 138,302 two-step solutions, harvested from simulations that use the distilled anchor net. On the years that fail the gate (5% wage gap and one percentage point of return), the hardest cases for the sharpened net, it cuts the median distance to that exact solve from 5.6% to 4.7% in wages and from 1.6 to 1.2 percentage points in returns.

In simulation the corrector is invoked only where it is needed. Each year is scored by one clearing evaluation; economies passing the gate keep their amortized heads, and only the failures are corrected.

Appendix E details the encoder input features and architecture, the clearing objective, the

training procedure, and the learned correction.

3.6 Posterior Estimation

With the CE and rollout agents’ decision rules and the anchor solver amortized, one equilibrium path for both countries is a sequence of cheap forward passes in the neural networks. However, because the equilibrium path involves nonlinear state transitions characterized by the neural networks and high-dimensional stochastic shocks, evaluating the likelihood of the observed data given parameter values is intractable. Drawing from that same likelihood, by contrast, is cheap, which is one simulated path per parameter draw. I therefore compute the posterior by simulation-based inference, which exploits this asymmetry: learning a direct map from observed outcomes to the parameter posterior and bypassing likelihood evaluation (Cranmer et al., 2020; Zammit-Mangion et al., 2024).

Let $\boldsymbol{\theta}$ collect the estimated parameters. Let $\mathbf{D}$ denote the observed data: for China (1983–2019) and Korea (1963–2019), the capital and output growth rates and the top-1%, top-10%, and bottom-50% pre-tax income and wealth shares. Wealth follows the WID net-worth definition, whose model counterpart is the saving (individual capital holding) $k_{t,a}$ of (39). Income is the pre-tax sum of labor and capital income, the agent’s gross labor earnings $w_{j,t} q_{t,a}$ plus the net capital return $r_{p,j,t} k_{t,a}$ on that wealth. The goal is the posterior $p(\boldsymbol{\theta} \mid \mathbf{D})$.

I approximate this posterior with a neural network $\Psi(\mathbf{D}; \boldsymbol{\omega})$ that takes any dataset $\mathbf{D}$ and returns the parameters of an approximate posterior $\hat{p}(\boldsymbol{\theta}; \Psi(\mathbf{D}; \boldsymbol{\omega}))$. I take $\hat{p}$ to be logit-Normal: each parameter is Gaussian on the logit of its prior support, with its own location and scale, and independent across parameters.¹⁷

The neural network weights $\boldsymbol{\omega}$ minimize the expected Kullback–Leibler divergence between the true and approximate posteriors, averaged over the marginal distribution of datasets,

$$\boldsymbol{\omega}^* = \arg \min_{\boldsymbol{\omega}} \int \left(\int p(\boldsymbol{\theta} \mid \mathbf{D}) \log \frac{p(\boldsymbol{\theta} \mid \mathbf{D})}{\hat{p}(\boldsymbol{\theta}; \Psi(\mathbf{D}; \boldsymbol{\omega}))} d\boldsymbol{\theta} \right) p(\mathbf{D}) d\mathbf{D}. \quad (68)$$

Intuitively, the inner integral measures how far the neural posterior sits from the true one at a given dataset, and the outer integral averages that distance over every dataset the model can generate, so $\boldsymbol{\omega}^*$ makes the neural posterior as close as possible to the true posterior on average.

Applying Bayes’ rule $p(\boldsymbol{\theta} \mid \mathbf{D}) = p(\mathbf{D} \mid \boldsymbol{\theta})p(\boldsymbol{\theta})/p(\mathbf{D})$ and dropping terms that do not depend

¹⁷Unlike a truncated Normal, a logit-Normal places no mass at the boundary of the support and needs no boundary correction. Because the logit scale is unbounded, the loss is a plain Gaussian negative log density and its gradients stay well conditioned, whereas a truncated Normal’s bound-dependent normalizing constant makes them fragile exactly where the posterior concentrates near a prior bound.

on ω simplifies the objective to

$$\omega^* = \arg \min_{\omega} - \int \left(\int p(\mathbf{D} \mid \boldsymbol{\theta}) \log \hat{p}(\boldsymbol{\theta}; \Psi(\mathbf{D}; \omega)) d\mathbf{D} \right) p(\boldsymbol{\theta}) d\boldsymbol{\theta}. \quad (69)$$

This expectation is first over the prior $p(\boldsymbol{\theta})$ and then over the likelihood $p(\mathbf{D} \mid \boldsymbol{\theta})$, which ancestral sampling estimates directly: drawing $\boldsymbol{\theta}_i \sim p(\boldsymbol{\theta})$ and simulating $\mathbf{D}_i \sim p(\mathbf{D} \mid \boldsymbol{\theta}_i)$ with the equilibrium simulator gives the Monte-Carlo objective

$$\omega^* \approx \arg \min_{\omega} - \sum_{i=1}^{N_{mc}} \log \hat{p}(\boldsymbol{\theta}_i; \Psi(\mathbf{D}_i; \omega)). \quad (70)$$

Each draw simulates one equilibrium path for both China and Korea. I set $N_{mc} = 50,000$.

I parameterize Ψ as a residual multilayer perceptron that maps the flattened, country-concatenated summary vector $\mathbf{D} \in \mathbb{R}^{752}$ to a logit-Normal approximate posterior over the 68 parameters, giving a location and scale for each (Figure 37), and minimize (70) by stochastic gradient descent.




Figure 37: Inference Network Dataflow

Notes: The left panel is the input $\mathbf{D}$: eight series per country, the capital and output growth rates and the top-1%, top-10%, and bottom-50% pre-tax income and wealth shares, over 37 years for China (1983–2019) and 57 years for Korea (1963–2019), flattened and concatenated into a vector in $\mathbb{R}^{752}$. The middle panel is a stylized rendering of the residual MLP $\Psi(\mathbf{D}; \omega)$; its exact architecture, a Dense(752 $\rightarrow$ 256, swish) projection, three residual blocks of width 256, and a linear Dense(256 $\rightarrow$ 136) head, is given in Appendix F. The 136 outputs are the location μ_p and scale σ_p of the logit-Normal approximate posterior of each of the 68 parameters; the right panel sketches one such marginal.

Passing the observed summaries $\mathbf{D}_{\text{obs}}$ of China and Korea through the trained Ψ returns the posterior for every parameter. Appendix F details the summary features, the network architecture and logit-Normal head, the training protocol, and the inference diagnostics.

3.7 Estimation Results

Passing the observed Chinese and Korean summaries through the trained inference network of Section 3.6 yields the joint posterior over the 68 estimated parameters. Table 1 sets each parameter’s prior beside its posterior mean, median, and 90% credible interval from 10,000

draws, together with its *contraction*: the posterior standard deviation of those draws divided by the prior standard deviation. In the prior column $\mathcal{TN}(\mu, \sigma^2; L, U)$ is a normal distribution with mean μ and standard deviation σ truncated to $[L, U]$, and $\text{TPois}(\lambda; L, U)$ a Poisson distribution with rate λ truncated to $\{L, \dots, U\}$.

Roughly a third of the parameters are sharpened by the data: 23 of 68 contract below 0.9 and six below 0.5. The macro time-series data identify the parameters that shift aggregate levels and trends directly, such as the mean discount factor μ_β , the mean risk aversion μ_γ , the labor-shock persistence η_ε , and the 1962 frontier level.

At the other end, the family preference dispersions such as $\ln \sigma_\beta$ and $\ln \sigma_\nu$ are not identified at all: the macro inequality series cannot separate one preference parameter's dispersion from another's. The forecast horizon Δ and the frontier growth rate g_f are likewise left at their priors.

Table 1: Prior and posterior distributions of the estimated parameters.

Param.	Description	Prior	Posterior				
		Distribution	Mean	Median	5%	95%	Contr.
<i>Frontier and TFP Adoption</i>							
g_f	Frontier TFP growth	$\mathcal{TN}(0.0055, 0.0015^2; 0.002, 0.01)$	0.005	0.005	0.003	0.008	1.01
$A_{f,1962}$	Frontier TFP level 1962	$\mathcal{TN}(28, 8.3^2; 3, 53)$	10.748	9.752	5.376	19.359	0.55
$\theta_{A,cn}$	CN technology ceiling	$\mathcal{TN}(0.8, 0.15^2; 0.5, 1)$	0.798	0.812	0.592	0.962	0.98
$\theta_{A,kr}$	KR technology ceiling	$\mathcal{TN}(0.8, 0.15^2; 0.7, 1)$	0.811	0.792	0.708	0.967	1.05
$\lambda_{cn,1982}$	CN initial TFP ratio	$\mathcal{TN}(0.2, 0.05^2; 0.1, 0.5)$	0.208	0.200	0.136	0.311	1.14
$\lambda_{kr,1962}$	KR initial TFP ratio	$\mathcal{TN}(0.3, 0.05^2; 0.1, 0.5)$	0.267	0.263	0.173	0.375	1.24
v_{max}	Maximum adoption speed	$\mathcal{TN}(0.12, 0.05^2; 0.05, 0.25)$	0.127	0.121	0.065	0.208	1.04
ω_h	High-school adoption weight	$\mathcal{TN}(1.5, 0.5^2; 1, 4)$	1.703	1.604	1.156	2.599	1.14
$\Delta\omega_c$	College weight increment	$\mathcal{TN}(2.0, 1.5^2; 0, 6)$	2.125	2.037	0.862	3.687	0.70
$\eta_{cn,u}$	CN TFP-shock persistence	Beta(4, 2)	0.634	0.644	0.390	0.838	0.77
$\eta_{kr,u}$	KR TFP-shock persistence	Beta(4, 2)	0.677	0.712	0.299	0.935	1.10
$\sigma_{cn,u}$	CN TFP-shock s.d.	$\mathcal{TN}(0.05, 0.025^2; 0, 0.1)$	0.051	0.051	0.009	0.091	1.17
$\sigma_{kr,u}$	KR TFP-shock s.d.	$\mathcal{TN}(0.05, 0.025^2; 0, 0.1)$	0.054	0.055	0.012	0.092	1.15
<i>Nested CES Production</i>							
$\alpha_{Y,cn}$	CN final-output CES share	$\mathcal{TN}(0.5, 0.15^2; 0.05, 0.95)$	0.648	0.664	0.411	0.834	0.88
$\alpha_{Y,kr}$	KR final-output CES share	$\mathcal{TN}(0.5, 0.15^2; 0.05, 0.95)$	0.540	0.543	0.314	0.757	0.91
$\log \sigma_{Y,cn}$	CN capital-labor elast.	$\mathcal{TN}(\log 0.6, 0.25^2; \log 0.3, \log 1.3)$	-0.324	-0.308	-0.726	0.019	0.93
$\log \sigma_{Y,kr}$	KR capital-labor elast.	$\mathcal{TN}(\log 0.6, 0.25^2; \log 0.3, \log 1.3)$	-0.075	-0.045	-0.405	0.149	0.70
$\alpha_{K,cn}$	CN private-capital share	$\mathcal{TN}(0.5, 0.2^2; 0.1, 0.95)$	0.447	0.423	0.162	0.805	1.12
$\alpha_{K,kr}$	KR private-capital share	$\mathcal{TN}(0.5, 0.2^2; 0.1, 0.95)$	0.467	0.462	0.252	0.703	0.76
$\sigma_{K,cn}$	CN capital aggregator elast.	$\mathcal{TN}(0.5, 0.2^2; 0.1, 0.95)$	0.437	0.404	0.141	0.833	1.22
$\sigma_{K,kr}$	KR capital aggregator elast.	$\mathcal{TN}(0.5, 0.2^2; 0.1, 0.95)$	0.466	0.454	0.213	0.760	0.94
δ	Depreciation rate	$\mathcal{TN}(0.05, 0.01^2; 0.01, 0.1)$	0.044	0.042	0.024	0.067	1.32

continued on next page

Table 1 – continued from previous page

Param.	Description	Prior	Posterior				
		Distribution	Mean	Median	5%	95%	Contr.
<i>Government Policy</i>							
$\Lambda_{cn,1}$	CN cons. floor (pre-2012)	$\mathcal{TN}(0.03, 0.02^2; 0, 0.2)$	0.035	0.030	0.008	0.082	1.33
$\Lambda_{cn,2}$	CN cons. floor (post-2012)	$\mathcal{TN}(0.1, 0.04^2; 0.02, 0.3)$	0.133	0.128	0.050	0.233	1.49
$\Lambda_{kr,1}$	KR cons. floor (pre-1988)	$\mathcal{TN}(0.03, 0.02^2; 0, 0.2)$	0.034	0.031	0.013	0.067	0.96
$\Lambda_{kr,2}$	KR cons. floor (post-1988)	$\mathcal{TN}(0.1, 0.04^2; 0.02, 0.3)$	0.090	0.083	0.043	0.160	0.97
$\tau_{k,cn}$	CN capital income tax	$\mathcal{TN}(0.10, 0.05^2; 0, 0.3)$	0.104	0.100	0.047	0.176	0.83
$\tau_{k,kr,1}$	KR capital tax (pre-1997)	$\mathcal{TN}(0.10, 0.05^2; 0, 0.3)$	0.091	0.088	0.048	0.145	0.63
$\tau_{k,kr,2}$	KR capital tax (post-1997)	$\mathcal{TN}(0.20, 0.05^2; 0, 0.3)$	0.181	0.185	0.094	0.257	1.06
$\tau_{m,kr}$	KR transfer tax rate	$\mathcal{TN}(0.3, 0.1^2; 0.1, 0.5)$	0.319	0.325	0.143	0.473	1.20
$\zeta_{cn,1}$	CN fiscal eff. (pre-break)	$\mathcal{TN}(0.7, 0.1^2; 0.5, 1)$	0.790	0.794	0.648	0.915	0.86
$\zeta_{cn,2}$	CN fiscal eff. (post-break)	$\mathcal{TN}(0.8, 0.1^2; 0.5, 1)$	0.800	0.805	0.677	0.909	0.76
$\zeta_{kr,1}$	KR fiscal eff. (pre-break)	$\mathcal{TN}(0.7, 0.1^2; 0.5, 1)$	0.786	0.789	0.656	0.903	0.79
$\zeta_{kr,2}$	KR fiscal eff. (post-break)	$\mathcal{TN}(0.8, 0.1^2; 0.5, 1)$	0.932	0.935	0.887	0.965	0.26
θ_κ	Students-to-teachers ratio	$\mathcal{TN}(20, 5^2; 10, 30)$	18.933	18.884	15.166	22.884	0.53
<i>Preferences and Family Heterogeneity</i>							
β_i	Family discount factor	$\mathcal{TN}(\mu_\beta, \sigma_\beta^2; 0.5, 1)$	—	—	—	—	—
μ_β	Mean discount factor	$\mathcal{TN}(0.95, 0.05^2; 0.8, 1)$	0.944	0.945	0.916	0.967	0.40
$\ln \sigma_\beta$	Log s.d. discount factor	$\mathcal{N}(-2.5, 0.5^2)$	-2.511	-2.519	-3.810	-1.217	1.58
γ_i	Family risk aversion	$\mathcal{TN}(\mu_\gamma, \sigma_\gamma^2; 0.5, 5)$	—	—	—	—	—
μ_γ	Mean risk aversion	$\mathcal{TN}(2.5, 0.5^2; 1.5, 3.5)$	2.994	3.035	2.511	3.333	0.57
$\ln \sigma_\gamma$	Log s.d. risk aversion	$\mathcal{N}(0, 0.2^2)$	-0.228	-0.250	-0.684	0.301	1.50
ψ_i	Family bequest weight	$\mathcal{TN}(\mu_\psi, \sigma_\psi^2; 0, 2)$	—	—	—	—	—
μ_ψ	Mean bequest weight	$\mathcal{TN}(1, 0.25^2; 0.5, 2)$	1.084	1.073	0.776	1.433	0.85
$\ln \sigma_\psi$	Log s.d. bequest weight	$\mathcal{N}(-0.7, 0.3^2)$	-0.618	-0.607	-1.415	0.148	1.58
$\log \chi_i$	Family disutility weight	$\mathcal{N}(\mu_{\log \chi}, \sigma_{\log \chi}^2)$	—	—	—	—	—
$\mu_{\log \chi}$	Mean log disutility weight	$\mathcal{N}(-1.92, 1^2)$	-3.851	-4.055	-5.763	-1.287	1.36
$\ln \sigma_{\log \chi}$	Log s.d. disutility wt.	$\mathcal{N}(-0.7, 0.5^2)$	-1.007	-1.036	-2.211	0.315	1.53
ν_i	Family disutility curvature	$\mathcal{TN}(\mu_\nu, \sigma_\nu^2; 0.2, 4)$	—	—	—	—	—
μ_ν	Mean disutility curvature	$\mathcal{TN}(1.0, 0.5^2; 0.2, 4)$	1.256	1.185	0.587	2.192	1.10
$\ln \sigma_\nu$	Log s.d. curvature	$\mathcal{N}(-1.0, 0.4^2)$	-1.065	-1.073	-2.076	-0.041	1.57
<i>Labor Productivity and Human Capital</i>							
η_ε	Labor-shock persistence	Beta(18, 2)	0.952	0.961	0.886	0.987	0.51
σ_ε	Labor-shock s.d.	$\mathcal{TN}(0.18, 0.06^2; 0.05, 0.4)$	0.207	0.204	0.117	0.308	1.02
Δh_h	High-school premium	$\mathcal{TN}(0.45, 0.12^2; 0.21, 0.75)$	0.510	0.511	0.436	0.579	0.39
Δh_c	College premium	$\mathcal{TN}(0.55, 0.15^2; 0.24, 0.9)$	0.566	0.564	0.344	0.792	1.00
δ_h	Human-capital depreciation	$\mathcal{TN}(0.02, 0.015^2; 0, 0.08)$	0.031	0.030	0.011	0.054	1.05
$\gamma_{h,0}$	Learning-by-doing scale	$\mathcal{TN}(0.06, 0.03^2; 0, 0.2)$	0.106	0.107	0.049	0.160	1.19
$\gamma_{h,1}$	Learning-age decline	$\mathcal{TN}(0.008, 0.004^2; 0, 1/70)$	0.009	0.009	0.003	0.013	0.98
α_h	Learning concavity	Beta(2, 2)	0.548	0.548	0.430	0.665	0.32
$\mathfrak{d}_h$	Learning smoothing	$\mathcal{TN}(0.1, 0.05^2; 0.01, 0.3)$	0.124	0.120	0.061	0.197	0.89
<i>Subjective Price Risk and CE Prediction</i>							


continued on next page

Table 1 – *continued from previous page*


Param.	Description	Prior	Posterior				
		Distribution	Mean	Median	5%	95%	Contr.
$\eta_{w,cn}$	CN wage-risk persistence	Beta(4, 2)	0.589	0.601	0.262	0.868	1.04
$\eta_{w,kr}$	KR wage-risk persistence	Beta(4, 2)	0.664	0.674	0.452	0.839	0.66
$\eta_{r,cn}$	CN return-risk persistence	Beta(4, 2)	0.612	0.630	0.277	0.882	1.04
$\eta_{r,kr}$	KR return-risk persistence	Beta(4, 2)	0.688	0.721	0.341	0.927	1.03
$\sigma_{w,cn}$	CN wage-risk s.d.	$\mathcal{TN}(0.05, 0.025^2; 0, 0.15)$	0.052	0.049	0.018	0.096	1.02
$\sigma_{w,kr}$	KR wage-risk s.d.	$\mathcal{TN}(0.05, 0.025^2; 0, 0.15)$	0.060	0.058	0.025	0.102	1.01
$\sigma_{r,cn}$	CN return-risk s.d.	$\mathcal{TN}(0.06, 0.03^2; 0, 0.18)$	0.063	0.059	0.022	0.116	1.02
$\sigma_{r,kr}$	KR return-risk s.d.	$\mathcal{TN}(0.06, 0.03^2; 0, 0.18)$	0.084	0.083	0.039	0.130	0.98
$\rho_{wr,cn}$	CN wage-return risk corr.	Unif(−0.75, 0.75)	0.180	0.181	0.080	0.274	0.14
$\rho_{wr,kr}$	KR wage-return risk corr.	Unif(−0.75, 0.75)	−0.005	−0.003	−0.660	0.653	0.99
Δ	Forecast horizon	TPois(5; 1, 10)	5.943	6.014	2.523	9.137	0.98
<i>Initial Aggregate State</i>							
$K_{s,cn}/K_{cn}$	CN initial state-cap. share	$\mathcal{TN}(0.6, 0.08^2; 0.5, 0.8)$	0.607	0.594	0.515	0.739	1.07
$K_{s,kr}/K_{kr}$	KR initial state-cap. share	$\mathcal{TN}(0.4, 0.1^2; 0.3, 0.7)$	0.368	0.365	0.337	0.407	0.27
$\mathcal{X}/K_p$	Initial endowment-capital ratio	$\mathcal{TN}(1.17, 0.1^2; 1, 1.5)$	1.163	1.143	1.029	1.362	1.14

Notes: $\mathcal{TN}(\mu, \sigma^2; L, U)$ is a normal distribution with mean μ and standard deviation σ truncated to $[L, U]$; TPois($\lambda; L, U$) is a Poisson distribution with rate λ truncated to the integers $\{L, \dots, U\}$.

Figure 38 checks the fit by posterior simulation. The model captures the long-run trajectories well: for China the growth rates and all the income shares lie inside the 90% band, with the median tracking the rise of the top shares, and for Korea the growth rates, wealth shares, and top income shares are likewise covered. Two tensions stand out. The model front-loads China’s wealth transition: the median top-10% wealth share rises above the data in the 1990s and the bottom-50% share falls below it, with the observed shares catching up to the model only after 2010. And the model holds Korea’s bottom-50% income share below the observed level over roughly 1985–2005: it cannot deliver as much income to the bottom half of Korean earners during the high-growth decades as the data show. However, such gaps reflect the model’s limited scope, which abstracts from some institutional changes that moved these distributions directly, such as the repeal of wage-suppression policies in 1980s Korea (Hong et al., 2025). They are not a failure of the inference, which the held-out calibration of Appendix F shows to be sound on average, apart from several parameters whose intervals are too tight.



(a) China (1982–2019)



(b) Korea (1962–2019)

Figure 38: Posterior Predictive Check

Notes: Each panel overlays the model's posterior-predictive distribution against the observed series. The shaded band is the 5–95% range and the solid line the median across posterior-draw re-simulations, in orange for China and salmon for Korea; the black line is the observed data. Panels are the capital and output growth rates and the top-1%, top-10%, and bottom-50% pre-tax income and wealth shares.

4 Counterfactual Experiments

With the model estimated, I ask what each engine of growth reviewed in [Section 1](#) contributed to growth and what it did to inequality, by shutting that engine off at its source and re-simulating the equilibrium. Every effect is reported as a distribution rather than a point estimate: each experiment runs at 500 posterior draws, and every counterfactual path is paired with its baseline, so their difference isolates the model-implied effect of the intervention on the causal graph. Unless otherwise noted, the numerical effects discussed below are posterior medians, and the accompanying tables report 90% credible intervals. For a clean comparison, I drop an abnormal tail of draws, those that clear markets far less well than their own baseline or end the sample at an implausible capital return.¹⁸

4.1 The TFP Catch-Up

The first counterfactual experiment asks what technological catch-up contributed to growth. I switch it off by setting the maximum adoption speed to zero, $v_{\max} = 0$, which freezes the TFP-to-ceiling ratio $\lambda_{j,t}$ at its estimated initial value in [\(4\)](#): the country still rides frontier growth, but never closes its gap to the ceiling.

[Table 2](#) reports the effect. On the output it is first order: by the end of the sample in 2019, per-capita output is 217% higher in China and 306% higher in Korea than in an economy that never closed its technology gap. However, TFP catch-up has no clear effect on wealth shares.

¹⁸A draw is flagged if its market-clearing loss [\(59\)](#) runs above its own baseline's by more than a factor of ten in the median over years, or by a factor of a hundred in any single year, or if its end-of-sample capital return exceeds 50% a year. A flagged draw is dropped from the baseline and from the counterfactual together, so the pairing stays intact.

Table 2: Effect of TFP Catch-up

	China			Korea		
	Median	5%	95%	Median	5%	95%
<i>Production block (percent higher with catch-up)</i>						
K_s (state capital)	169.7	36.7	778.9	300.0	70.2	1480.4
K_p (private capital)	159.1	23.4	906.4	271.2	45.6	1427.8
working time	-15.6	-45.0	-0.8	-18.5	-48.5	-0.1
H (human capital)	-10.5	-31.3	-0.8	-18.1	-40.7	-4.8
A (TFP)	134.9	61.3	291.4	195.0	96.5	375.8
K (capital aggregate)	165.9	32.1	794.0	279.8	56.4	1396.2
L (effective labor)	-22.9	-56.5	-2.1	-28.8	-62.2	-7.6
Y (output)	216.5	56.6	922.2	305.8	84.7	1288.4
<i>Income and wealth shares (percentage points higher with catch-up)</i>						
Income, top 1%	0.6	-11.4	13.9	1.8	-7.8	17.5
Income, top 10%	1.3	-8.6	11.8	4.3	-3.4	16.1
Income, bottom 50%	-0.5	-5.0	3.0	-2.2	-7.2	1.6
Wealth, top 1%	0.1	-39.9	26.6	0.7	-34.5	34.0
Wealth, top 10%	-3.5	-29.4	14.2	-0.4	-22.3	17.7
Wealth, bottom 50%	0.6	-3.0	7.1	0.4	-4.1	6.7

Notes: Paired effect of catch-up in 2019, baseline against the $v_{\max} = 0$ counterfactual, over the 500 posterior draws. Production-block rows give the percent by which catch-up raises the level, $100(x_{base}/x_{cf} - 1)$; share rows the percentage-point gap, $100(s_{base} - s_{cf})$.

Figure 39 traces how the shock propagates through the distributions of nodes in the production block. Catch-up accelerates total factor productivity A by 135% in China and 195% in Korea, and deepens private and state capital by more still. The labor block moves the other way and by far less: working time falls by about a sixth, since higher productivity buys leisure, and human capital follows it down mildly. Catch-up therefore reached Chinese and Korean growth through productivity and the capital deepening it induced, not through labor or human-capital accumulation.




Figure 39: Propagation of the TFP Catch-up Counterfactual

Notes: The production network of (1), each node drawn as a kernel density across posterior draws in 2019, China on the upper sub-panel and Korea on the lower: the baseline posterior simulation (blue) and the no-catch-up counterfactual $v_{\max} = 0$ (red). Each sub-panel is standardized by its own baseline distribution and carries its own axis, so the horizontal axis is in that country's baseline standard deviations and 0 marks its baseline mean.

On wealth inequality, by contrast, TFP catch-up has no clear effect on wealth shares, as Figure 40 shows. It enters at productivity and reaches households only through aggregate prices, and the two prices it moves pull in opposite directions: the wage rises more than fivefold by 2019 in China, favoring families who live on labor income, while the private return rises by up to ten percentage points at the height of the transition, favoring the capital-rich. The factor-share channel that might break the tie is weak, since the estimated capital-labor elasticity of the CES production function is not far below Cobb–Douglas, a median σ_Y of 0.7 for China and

0.96 for Korea. Nor does the labor-supply response tilt the balance: catch-up lowers working time almost uniformly across education groups, so it cannot do much to amplify labor-income inequality.¹⁹




Figure 40: Wealth Shares With and Without Catch-up

Notes: The top-1%, top-10%, and bottom-50% wealth shares through time. Each line is the median across the gated posterior draws, with a 90% credible band.

4.2 The Education Expansion

Both countries expanded schooling substantially over the sample window. This experiment asks what that expansion contributed, by freezing the enrollment shares: high-school and college entry rates hold at their 1982 level in China and at their 1962 level in Korea.

Table 3 reports the effect as the contribution of the expansion at two horizons: 2019, the end of the estimation sample, and 2040, far enough out that the cohorts schooled after the expansion make up a substantial share of the workforce, since China's large-scale college expansion only began in 1999. Output is higher with the expansion at both horizons, by 3.6% in China and 16.6% in Korea in 2019, widening to 25.7% and 39.9% by 2040.

¹⁹The fall in working time is 14.5, 13.6, and 15.8 percent for the below-high-school, high-school, and college groups in China, and 21.2, 21.5, and 20.3 percent in Korea, from an education-conditional re-simulation.

Table 3: Contribution of the Education Expansion

	China						Korea					
	2019			2040			2019			2040		
	Median	5%	95%	Median	5%	95%	Median	5%	95%	Median	5%	95%
<i>Production block (percent change)</i>												
$\tilde{H}$ (human-capital intensity)	24.3	7.6	64.8	127.0	59.5	248.7	85.6	38.4	164.9	191.8	110.0	301.1
working time	-7.7	-15.2	3.1	-10.6	-29.6	12.7	-9.7	-24.6	1.4	-10.0	-30.7	0.8
A (TFP)	2.8	0.9	7.8	12.9	4.5	30.1	7.5	2.0	17.2	12.1	2.0	26.7
L (effective labor)	-0.1	-5.8	9.6	7.8	-8.9	34.4	9.0	-5.3	27.9	23.7	-1.2	51.8
K (capital aggregate)	1.5	-12.8	15.1	18.6	-13.6	63.5	8.4	-15.6	39.6	30.0	-17.4	88.8
Y (output)	3.6	-4.3	16.1	25.7	3.5	68.1	16.6	-3.0	45.2	39.9	5.5	85.8
<i>Income and wealth shares (percentage points)</i>												
Income, top 1%	0.1	-1.3	2.2	0.5	-3.9	5.8	0.3	-3.0	5.4	0.8	-3.1	7.5
Income, top 10%	0.5	-1.5	3.5	2.0	-3.4	6.7	0.9	-2.8	4.8	2.2	-3.2	7.1
Income, bottom 50%	-1.1	-3.6	0.2	-1.9	-5.7	0.2	-1.1	-4.2	0.4	-1.6	-4.5	0.3
Wealth, top 1%	-0.1	-5.8	7.1	0.2	-19.4	10.6	0.7	-10.4	10.8	2.0	-9.9	17.3
Wealth, top 10%	-0.2	-5.7	5.9	0.1	-14.2	9.6	0.6	-8.7	7.7	1.9	-10.3	12.5
Wealth, bottom 50%	-0.0	-2.2	1.7	-0.1	-3.1	3.8	-0.0	-1.9	3.6	-0.1	-2.5	3.5

Notes: Paired contribution of the observed education expansion, baseline against frozen enrollment, over the 500 gated posterior draws per country. Production-block rows give the percent change, $100(x_{base}/x_{cf} - 1)$; share rows the percentage-point change, $100(s_{base} - s_{cf})$. Years after the estimation sample are a projection in which the data tables, enrollment included, hold at their terminal values.

Figure 41 shows the same contributions as they travel through the model's causal graph. The expansion raises the human-capital intensity by 24% in China and 86% in Korea by 2019, and by 127% and 192% by 2040. That intensity gain passes into productivity: TFP is 2.8% and 7.5% higher in 2019, and 12.9% and 12.1% higher by 2040. Effective labor rises as the human-capital gain outgrows the shorter working time, reaching 7.8% in China and 23.7% in Korea by 2040, though in China the two still offset each other at 2019. The capital aggregate also rises, consistent with higher productivity raising the return and inducing saving.




Figure 41: Propagation of the Education Expansion

Notes: The treated node (left) is the enrollment rates. Every other node is a kernel density across posterior draws of the per-draw paired difference $\log(\text{expansion}) - \log(\text{frozen})$, the contribution of the observed expansion to that node in log points, where 0.1 is about 10% higher, at 2019 and at 2040; the dashed grey rule is zero and the coloured dashed lines are the medians. Years after 2019 are a projection in which enrollment holds at its terminal value.

Why is the 2019 effect so much smaller than the 2040 one? The answer lies in the adoption speed $v_{j,t} = v_{\max} \mathcal{H}_{j,t}$ of (5). The early gains in the human-capital density do not come from schooling. Far from the ceiling the low-hanging fruit is abundant, since (4) makes the proportional gain from a given speed largest when λ is smallest, so even the modest speed of human-capital accumulation sustained by an unschooled workforce produces rapid catch-up. What is left as λ approaches one is the stretch that only high school or college graduates with more human capital can reach.

Figure 42 bears this out: education buys the later stretch of the catch-up, not the first, which is why an experiment that stops at 2019 understates what the expansion is for. The gap between the two countries at 2019 is the same mechanism seen from the other side. Korea

expanded early, and from a position already closer to the technology ladder, whereas China's large-scale expansion arrived roughly two decades later. That is why the expansion's 2019 contribution to output is 16.6% in Korea against 3.6% in China.




Figure 42: The Catch-up Ratio With and Without the Education Expansion

Notes: The catch-up ratio $\lambda_{j,t}$ of (4), China (left) and Korea (right). The upper row plots the posterior median of λ under the observed enrollment (blue) and under frozen enrollment (red); the lower row plots the per-draw gap between them.

The distributional footprint is on income, and it is modest, because a broad expansion redistributes to the upper middle rather than to the very top: the top percentile's share moves by only 1 to 5 percent of its own level, while the bottom half, the households the expansion leaves behind, gives up 1.1 points in each, out of shares of 13 and 10 points. The education expansion has no clear effect on wealth shares, consistent with dilution by the dispersion in preferences and by the idiosyncratic labor-productivity shocks that drive wealth accumulation.

4.3 State Investment and Fiscal Efficiency

One of the most salient features of both economies is that the government accumulates a great deal of state capital, and the fiscal efficiency $\zeta_{j,t}$ of (22) governs how much fiscal revenue is converted into real expenditure, henceforth state investment. This section asks what each contributes to growth and to distribution.

The Impact of State-Owned Capital To study the impact of state capital, I run a “privatization” experiment: the government hands households a share π of each year’s state investment, credited uniformly per capita to their private capital, as a citizens’ dividend would be.

Table 4 reports the effect. Redirecting $\pi = 10\%$ of state investment lowers state capital by 31% in China and 44% in Korea by 2019, and output by 13.9% and 17.2%; doubling the share to 20% takes state capital down 49% and 64%, and output down 22.0% and 26.2%.

Table 4: Redirecting State Investment to Households

	China						Korea					
	$\pi = 10\%$			$\pi = 20\%$			$\pi = 10\%$			$\pi = 20\%$		
	Median	5%	95%	Median	5%	95%	Median	5%	95%	Median	5%	95%
<i>Production block (percent change)</i>												
K_s (state capital)	-30.6	-55.8	-10.1	-49.1	-73.0	-27.5	-43.9	-65.8	-21.8	-63.8	-81.3	-42.3
K_p (private capital)	-6.6	-49.8	44.3	-6.5	-60.1	59.5	-3.1	-53.9	60.3	-3.8	-61.8	87.0
K (capital aggregate)	-17.3	-50.6	18.5	-26.0	-65.0	14.8	-27.1	-60.2	3.8	-39.7	-74.5	1.2
A (TFP)	-0.5	-3.2	0.3	-0.6	-4.7	0.8	-0.8	-4.8	0.2	-0.9	-5.2	0.6
L (effective labor)	-6.7	-29.2	2.5	-8.0	-33.6	6.5	-9.8	-33.5	8.0	-11.9	-37.6	11.2
working time	-5.9	-28.2	3.9	-7.0	-32.0	9.2	-9.6	-36.8	10.6	-11.3	-40.4	11.4
w (wage)	-8.6	-31.0	8.3	-17.9	-49.4	1.0	-8.0	-36.7	8.7	-16.6	-48.2	2.3
H (human capital)	-3.1	-16.8	0.5	-4.1	-21.2	1.1	-7.0	-22.3	0.1	-7.8	-26.2	1.7
Y (output)	-13.9	-44.9	-0.9	-22.0	-57.9	-4.8	-17.2	-49.3	3.9	-26.2	-61.5	1.2
<i>Income and wealth shares (percentage points)</i>												
Income, top 1%	-0.3	-16.7	1.9	-0.2	-24.8	2.4	-0.4	-23.4	2.6	-0.1	-29.5	3.7
Income, top 10%	0.1	-12.9	4.9	0.6	-18.4	7.0	-0.2	-16.6	6.1	1.1	-22.4	9.0
Income, bottom 50%	-0.2	-3.6	5.3	-0.9	-6.1	7.1	-0.2	-4.3	6.9	-1.4	-6.7	8.6
Wealth, top 1%	-11.8	-53.1	-0.3	-16.1	-68.0	-0.8	-9.9	-58.1	-0.6	-12.7	-73.1	-0.8
Wealth, top 10%	-15.0	-46.0	-1.0	-23.4	-57.7	-2.9	-17.6	-42.5	-2.3	-24.3	-58.5	-3.5
Wealth, bottom 50%	6.0	1.6	17.3	10.0	3.1	22.4	8.5	2.3	15.9	11.8	3.7	21.4

Notes: Effect of the intervention in 2019 over 500 jointly gated posterior draws, counterfactual minus baseline. Percent-change rows are $100(x_{cf}/x_{base} - 1)$; percentage-point rows are $100(s_{cf} - s_{base})$. A draw is dropped from both rungs at once, so the columns are a dose-response curve on one sample per country. The median and the 5–95% band are reported rather than the mean, which a few draws in weakly identified corners dominate.

Figure 43 traces where the loss comes from. Effective labor falls more than private capital. The transfer is associated with declines in working time of 5.9 to 11.3%. Shorter working time then feeds back through learning by doing, and human capital falls by 3.1 to 7.8%. Effective labor carries both losses into output, with declines of 6.7 to 11.9%.

On the distribution, redirecting 20% of state investment lowers the top-10% wealth share by 23.4 percentage points in China and 24.3 in Korea, and raises the bottom half’s share by 10.0 and 11.8 points. Figure 44 shows the process: a per-capita credit paid every year for four decades compounds into the lower tail.

The experiment is at least suggestive of a growth-inequality trade-off in a developmental state: relative to the rebate, state-led capital accumulation supports output but leaves less wealth at the bottom.




Figure 43: Propagation of the State-Investment Rebate

Notes: Each node is a kernel density across the posterior draws of the per-draw paired difference $\log(\text{counterfactual}) - \log(\text{baseline})$ in 2019.




Figure 44: Wealth Distribution Under the Rebate

Notes: The right panel draws each a median with a 90% credible band.

The Impact of Fiscal Efficiency The second experiment cuts fiscal efficiency by 10%, 20%, or 50%. The lower efficiency makes the resources leave the economy entirely rather than reaching households. Table 5 reports the result. A 10% cut shrinks state capital by 28% against 31% in China, yet output falls by 5.6% against 13.9%. And unlike the “privatization” experiment, the 10% and 20% cuts have small effects on the income and wealth shares, with no sign clearly determined.

Table 5: Lower Government Fiscal Efficiency

China									
	$\Delta\zeta/\zeta = -10\%$			$\Delta\zeta/\zeta = -20\%$			$\Delta\zeta/\zeta = -50\%$		
	Median	5%	95%	Median	5%	95%	Median	5%	95%
<i>Production block (percent change)</i>									
K_s (state capital)	-27.9	-43.9	-16.6	-46.7	-67.0	-32.6	-79.9	-95.2	-65.3
K_p (private capital)	-11.8	-38.5	15.6	-22.8	-59.4	9.4	-54.7	-91.4	-1.2
K (capital aggregate)	-18.9	-41.6	-2.2	-33.3	-63.1	-10.4	-66.3	-93.1	-32.9
A (TFP)	0.5	0.0	2.2	1.0	0.1	4.3	2.9	0.3	11.9
L (effective labor)	3.2	-5.4	18.4	6.3	-1.0	36.9	21.1	2.1	107.5
working time	1.9	-8.0	14.4	3.8	-5.2	25.1	11.6	0.6	71.9
w (wage)	-12.7	-34.5	-1.9	-24.0	-57.3	-7.2	-58.6	-92.1	-28.2
H (human capital)	2.0	0.1	8.2	3.6	0.3	15.7	10.7	1.2	45.9
Y (output)	-5.6	-25.8	2.5	-11.5	-44.0	0.4	-32.3	-83.0	-3.5
<i>Income and wealth shares (percentage points)</i>									
Income, top 1%	-0.1	-4.8	1.3	-0.3	-9.2	2.1	-0.7	-16.3	6.9
Income, top 10%	-0.4	-5.2	1.6	-0.8	-7.8	1.6	-2.0	-16.1	5.3
Income, bottom 50%	0.2	-0.9	2.2	0.4	-0.8	3.3	0.9	-2.2	5.4
Wealth, top 1%	-0.3	-12.9	8.0	-0.7	-23.6	9.3	-1.2	-41.1	27.3
Wealth, top 10%	-0.2	-11.9	6.3	-0.8	-16.9	8.2	-1.7	-30.6	17.4
Wealth, bottom 50%	0.0	-2.4	3.4	0.0	-2.4	4.6	0.1	-3.3	8.2
Korea									
	$\Delta\zeta/\zeta = -10\%$			$\Delta\zeta/\zeta = -20\%$			$\Delta\zeta/\zeta = -50\%$		
	Median	5%	95%	Median	5%	95%	Median	5%	95%
<i>Production block (percent change)</i>									
K_s (state capital)	-42.2	-56.0	-30.3	-63.0	-77.1	-50.8	-89.3	-96.5	-80.0
K_p (private capital)	-15.7	-43.6	5.2	-27.6	-60.0	0.8	-56.6	-88.9	-5.5
K (capital aggregate)	-31.8	-50.2	-16.6	-48.9	-70.7	-29.7	-76.0	-93.0	-50.8
A (TFP)	0.5	0.0	1.7	1.0	0.1	3.2	2.6	0.4	8.3
L (effective labor)	4.1	-5.6	18.8	8.4	-0.2	33.0	22.0	4.3	92.0
working time	2.4	-8.2	18.5	4.7	-3.5	27.8	12.5	-0.3	62.3
w (wage)	-14.7	-37.4	-2.9	-26.2	-55.4	-8.7	-55.8	-85.3	-23.8
H (human capital)	2.4	-0.3	8.0	5.0	0.7	15.0	14.0	3.4	38.2
Y (output)	-8.9	-33.1	2.1	-16.6	-51.7	0.2	-38.4	-81.4	-4.5
<i>Income and wealth shares (percentage points)</i>									
Income, top 1%	-0.5	-6.8	0.9	-1.0	-13.2	0.3	-2.0	-26.8	0.5
Income, top 10%	-1.0	-5.6	1.3	-1.9	-10.0	0.7	-3.9	-22.0	0.2
Income, bottom 50%	0.4	-0.7	2.1	0.8	-0.6	3.5	1.7	-0.9	7.2
Wealth, top 1%	-1.0	-12.1	5.1	-2.1	-20.6	4.0	-4.9	-45.1	4.3
Wealth, top 10%	-1.4	-8.6	4.1	-2.8	-14.5	5.6	-6.3	-31.7	8.2
Wealth, bottom 50%	0.2	-1.5	2.2	0.3	-2.2	3.4	0.9	-4.2	7.1

Notes: Effect of the intervention in 2019 over 500 gated posterior draws, counterfactual minus baseline. Percent-change rows are $100(x_{cf}/x_{base} - 1)$; percentage-point rows are $100(s_{cf} - s_{base})$. A draw is dropped from every rung at once, so the columns are a dose-response curve on one sample. The median and the 5–95% band are reported rather than the mean, which a few draws in weakly identified corners dominate.

Figure 45 shows the labor feedback that cushions the output drop. Households lose on both of their income margins at once. The two capitals of (7) are gross complements at the estimated

elasticity of substitution, a median σ_K of 0.40 in China and 0.45 in Korea, so private capital falls with the state stock, by 11.8 to 54.7% in China. Capital income falls by 13.2 to 59.3% in China and 25.7 to 79.0% in Korea, while the wage falls 13% to 59% in China. Working time rises by 2% to 12%; that response feeds learning by doing into human capital and, through (5), into the adoption speed, leaving productivity 0.5 to 2.9% higher under a less efficient government. This response offsets part of the capital loss.




Figure 45: Propagation of Lower Fiscal Efficiency

Notes: Each node is a kernel density across the posterior draws of the per-draw paired difference $\log(\text{counterfactual}) - \log(\text{baseline})$ in 2019; the dashed grey rule is zero and the coloured dashed lines are the medians. The treated node on the left shows the estimated ζ path of (23), solid, against each cut, dashed. Fiscal efficiency reaches working time directly, and not only through prices, because it enters agents' information set and their predicted state capital.

At the two shallow cuts of efficiency, the change in every wealth share is less than 3 percentage points in either country and no sign is determined; even halving efficiency changes the top-10% share by only 6.3 points in Korea and 1.7 in China. The cut shrinks the state's capital without ever passing through a household budget, so it reaches the distribution only through aggregate prices. That is the main difference from the rebate in the “privatization” experiment, which moves the same aggregate through household wealth and produces shifts in the shares that are an order of magnitude larger.

Taken together, the results indicate that state capital in these economies is productive. Redirecting the resources to households reduces wealth concentration but has a large cost in

output, with the decline in effective labor exceeding that in private capital. The efficiency with which fiscal revenue becomes expenditure also has large effects on output. This finding is consistent with the literature on state capacity, which treats it as the key to East Asia’s success (Johnson, 1982; Amsden, 1989; Kohli, 2004).

4.4 The Social Safety Net

This experiment asks what a social safety net does to growth and to inequality. The consumption floor of (14) is a ratio to the wage that takes one value before each country’s policy break and another after. I raise both regimes to a common ratio, first 0.2 and then 0.3, and compare the result with the posterior estimates. The posterior medians of Λ are 0.030 before the break in China and 0.031 in Korea, rising to 0.128 and 0.083 after it. Table 6 reports the aggregates. A floor of 0.2 lowers output by 3.2% in China and 7.6% in Korea; raising it to 0.3 lowers output by 10.0% and 19.3%.

Table 6: Raising the Consumption Floor

	China						Korea					
	$\underline{x}/w = 0.2$			$\underline{x}/w = 0.3$			$\underline{x}/w = 0.2$			$\underline{x}/w = 0.3$		
	Median	5%	95%	Median	5%	95%	Median	5%	95%	Median	5%	95%
<i>Production block (percent change)</i>												
K_s (state capital)	-2.1	-16.9	17.1	-9.2	-36.4	11.3	-7.8	-27.6	5.5	-21.1	-51.9	0.3
K_p (private capital)	2.1	-27.5	66.0	-7.8	-47.3	57.3	-6.6	-36.6	32.5	-22.5	-60.0	27.6
K (capital aggregate)	0.6	-20.8	42.2	-7.9	-40.4	33.5	-7.0	-29.2	14.1	-23.1	-53.6	7.1
A (TFP)	-0.3	-1.7	0.5	-0.8	-3.5	0.5	-0.5	-2.4	0.2	-1.0	-4.6	0.4
L (effective labor)	-4.5	-22.9	3.7	-10.7	-33.6	1.6	-7.9	-24.7	3.5	-16.4	-41.8	4.0
working time	-2.6	-19.8	5.5	-7.2	-30.4	5.4	-5.7	-23.8	6.6	-12.9	-38.9	8.4
H (human capital)	-2.0	-9.9	1.3	-4.1	-18.4	0.6	-4.1	-13.7	0.4	-9.0	-27.3	0.3
Y (output)	-3.2	-20.0	8.0	-10.0	-34.8	4.8	-7.6	-24.4	3.6	-19.3	-46.7	-0.0
<i>Wealth shares (percentage points)</i>												
Wealth, top 1%	-3.5	-29.2	6.3	-2.5	-36.7	18.8	-1.0	-21.2	7.9	0.1	-26.1	26.5
Wealth, top 10%	-4.0	-24.8	5.8	-3.5	-29.9	14.7	-2.2	-19.4	8.3	0.3	-22.2	24.9
Wealth, bottom 50%	0.5	-2.5	5.9	0.3	-3.4	7.3	0.5	-2.9	5.7	-0.0	-6.1	6.7
<i>Consumption levels (percent change)</i>												
Consumption at the 1st pct.	36.3	-35.6	289.8	84.4	-20.3	535.4	97.2	-17.5	341.4	182.0	-9.9	724.4
Consumption at the 5th pct.	17.9	-29.5	204.3	56.0	-26.8	343.9	61.2	-27.8	262.2	114.0	-25.0	418.0
Consumption at the median	-2.1	-27.2	34.0	-7.5	-39.6	39.8	-11.1	-40.1	27.0	-21.4	-53.8	20.5

Notes: Effect of raising the consumption floor in 2019 over 500 jointly gated posterior draws, counterfactual minus baseline. Percent-change rows are $100(x_{cf}/x_{base} - 1)$; percentage-point rows are $100(s_{cf} - s_{base})$. A draw is dropped from both rungs at once, so the columns rest on one sample per country. The median and the 5–95% band are reported rather than the mean, which a few draws in weakly identified corners dominate.

Figure 46 traces the propagation of that cost. State capital falls by 2.1 and 9.2% in China during 1982 to 2019, and 7.8 and 21.1% in Korea during 1962 to 2019. Working time also falls, by 2.6 to 7.2% in China and 5.7 to 12.9% in Korea, consistent with low-wealth agents choosing more leisure; private capital falls at the larger dose.




Figure 46: Propagation of a More Generous Consumption Floor

Notes: China on the upper sub-panel of each node and Korea on the lower. Each node is a kernel density across the gated posterior draws of the per-draw paired difference $\log(\text{counterfactual}) - \log(\text{baseline})$ in 2019, in log points, so a negative value means the higher floor lowered that node and -0.1 is about 10% lower; the dashed grey rule is zero and the coloured dashed lines are the medians. The treated node on the left shows the estimated Λ_j path of (14), solid, against each target, dashed.

Figure 47 shows the incidence of the policy. By 2019 a floor of 0.2 raises the amount consumed at the first percentile by 36% in China and 97% in Korea, and a floor of 0.3 raises it by 84% and 182%; the fifth percentile gains 18% to 114%. The effect at the tenth percentile is close to zero in 2019, and the median household consumes 2.1 to 21.4% less. The figure also shows the gain decaying at each country's policy break, where the estimated floor steps up on its own and leaves less room to raise it.

The smaller increase produces substantial consumption gains at the bottom relative to its output cost, although the credible intervals are wide. Comparing medians, each percent of output forgone corresponds to about 11% more consumption at China's first percentile and 13% at Korea's. A larger floor is a different proposition: raising it to 0.3 costs 10.0% of China's output in 2019 and 19.3% of Korea's. This trade-off is consistent with accounts of why developmental states tend to hold welfare provision down in the early stages of development (Kohli, 2004).




Figure 47: The Impact of Safety Net on Consumption

Notes: The per-draw paired difference $\log(\text{counterfactual}) - \log(\text{baseline})$ in four quantiles of the consumption cross-section. Each line is the median across the jointly gated posterior draws; the dashed grey rule is zero.

4.5 The Fertility Decline

Both countries experienced sharp fertility declines over the sample period. This experiment assesses their effects by freezing each country's birth rate: China's at its 1978 level, corresponding to the last cohort conceived before the one-child policy, and Korea's at its 1983 level, when its total fertility rate first fell below 2.1. Agents enter the model at age 15, so those cohorts arrive in 1993 and 1998, and the frozen rates apply from then on. Table 7 reports the effects in 2019. The resulting per-capita gains mainly reflect smaller populations, which fell by 4.9% in China and 12.2% in Korea, rather than higher output, which changed by 0.6% and -3.6% , respectively. This is the demographic dividend as the model sees it: the share of the population past schooling age rises by 3.0 percentage points in China and 5.0 in Korea.

The fertility decline has no clear positive effect on output in either economy in 2019: output changes by 0.6% in China and -3.6% in Korea, and both 90% bands cross zero. The results therefore attribute the per-capita gain to a smaller population rather than greater scale, and the cohorts it removed would have been the workers of the decades after 2019.

Carrying simulations for both countries forward makes that price concrete. Holding the entrant rate and life expectancy at their 2019 values and letting the age distribution follow, the decline lowers output by 6.6% in China by 2030 and by 8.9% in Korea by 2025. Output per

capita stays higher throughout, by 9.2% and 11.9%, so both economies end up richer per capita and materially smaller, and the gap in the level widens with every year the missing cohorts would have been at work.

Table 7: Contribution of the Fertility Decline

	China			Korea		
	Median	5%	95%	Median	5%	95%
<i>Population and production (percent change)</i>						
P (population)	-4.9	-4.9	-4.9	-12.2	-12.2	-12.2
P_w (working-age population)	-1.7	-2.2	-1.2	-7.1	-7.4	-6.9
K (capital aggregate)	1.1	-10.0	13.4	-1.5	-12.5	7.4
A (TFP)	0.2	-0.0	0.3	-0.2	-0.5	-0.0
L (effective labor)	0.1	-7.3	8.3	-4.3	-13.1	3.4
working time	-2.4	-9.7	6.9	-8.0	-15.9	1.6
H (human capital)	0.2	-2.1	1.3	-4.3	-9.4	-1.0
Y (output)	0.6	-5.0	7.0	-3.6	-10.8	2.5
Y/P (output per capita)	5.7	-0.1	12.5	9.8	1.6	16.8
<i>Age structure and shares (percentage points)</i>						
Working-age share P_w/P	3.0	2.6	3.4	5.0	4.8	5.2
Income, top 1%	-0.3	-1.8	1.2	-0.5	-2.0	1.4
Income, top 10%	-0.8	-2.4	1.5	-1.1	-3.1	1.0
Income, bottom 50%	1.0	-0.1	2.0	1.1	0.1	2.5
Wealth, top 1%	-0.4	-4.2	6.5	-0.7	-3.6	4.2
Wealth, top 10%	-0.8	-5.8	4.7	-1.1	-4.7	3.1
Wealth, bottom 50%	0.2	-0.7	1.7	0.4	-0.4	2.5

Notes: Effect of the fertility decline in 2019 over 500 gated posterior draws, observed minus counterfactual, so a positive entry means the change raised that row. Percent-change rows are $100(x_{cf}/x_{base} - 1)$; percentage-point rows are $100(s_{cf} - s_{base})$.

The distributional footprint is small. The decline lowered the top 10% income share by 0.8 percentage points in China and 1.1 in Korea and raised the bottom half's by 1.0 and 1.1, while no wealth share has a band excluding zero. This largely reflects composition rather than redistribution: the cohorts not born would have been concentrated near the bottom of the income distribution, as students without labor income and as first-year workers with the least accumulated human capital, so the effect raises income at the fifth percentile by 260% in China against 2.5% at the ninety-ninth percentile, while the labor share of income moves 0.07 percentage points. As with the education expansion, an income effect of this size is diluted by the dispersion in preferences and by the idiosyncratic shocks that drive wealth accumulation.

4.6 Other Sources of Inequality

The five preceding experiments each shut off an engine of growth, and several of them moved the income and wealth distribution as well. Those are policy levers. This section asks what the model's primitives contribute, the preferences, risk, technology and tax rates that were estimated rather than chosen, by shutting off one source at a time while holding every other parameter at its estimated value. Seven sources are treated this way:

1. the dispersion across families of the discount factor β_i , frozen to zero, so that every family holds the mean patience of its own economy;
2. the dispersion across families of risk aversion γ_i , frozen to zero in the same way;
3. the idiosyncratic labor-productivity shock of (27), set to zero;
4. the aggregate productivity shock of (6), set to zero;
5. the labor-capital substitution elasticity of (1), moved for China to the Cobb-Douglas case $\sigma_{Y,\text{cn}} = 1$;
6. the capital tax, set to zero in both countries;
7. the inheritance tax, set to zero for Korea.

Comparing the posterior simulations with those generated by each source off gives the distribution of that source’s marginal contribution to inequality. Table 8 reports the median of each contribution in percentage points: a positive entry on a top share means the source raises wealth or income concentration.

Table 8: Shutting Off One Source of Inequality at a Time

	Income share			Wealth share		
	Top 1%	Top 10%	Bot 50%	Top 1%	Top 10%	Bot 50%
China						
discount-factor dispersion	-0.00	0.10	-0.07	0.71	1.58	-0.42
risk-aversion dispersion	1.22	3.07	-1.75	3.35	3.53	-0.52
labor productivity shock	4.08	9.59	-5.13	4.74	6.24	-1.20
aggregate shock	0.01	0.02	-0.01	-0.07	-0.07	0.00
labor-capital elasticity	-4.36	-4.74	1.31	-5.12	-3.19	0.27
capital tax	0.09	0.03	0.01	0.65	0.43	-0.02
Korea						
discount-factor dispersion	0.14	0.37	-0.21	0.86	1.49	-0.26
risk-aversion dispersion	0.86	2.50	-1.78	2.36	3.47	-0.72
labor productivity shock	5.62	10.51	-3.96	7.00	8.57	-1.29
aggregate shock	0.00	0.00	-0.00	0.01	-0.03	0.00
capital tax	0.05	-0.20	0.14	0.83	0.32	-0.00
bequest tax	0.09	0.16	0.01	0.68	0.72	-0.01

Notes: Contribution of each source to inequality in 2019, baseline minus the counterfactual, in percentage points, over 500 posterior draws. A positive entry on a top share means the source raises concentration.

Figure 48 compares each channel’s impact to the wealth distribution and shows the ranking with its uncertainty.²⁰

²⁰Every band crosses zero because in some samples dispersion lowers wealth concentration rather than raising it. Preferences are drawn independently of everything else that makes a family rich, so in a panel where a family already wealthy for other reasons may happen to draw a low discount factor or a high risk aversion. Dispersion then holds that family back, and removing it lets the family accumulate freely, which concentrates wealth further. Whether the draw falls that way is sampling luck, so it varies across draws while the median recovers the systematic effect.




Figure 48: Sources of Inequality, Ranked

Notes: China on the upper row and Korea on the lower, each with its own sources and its own ordering. Each dot is the median contribution of that source, baseline minus the counterfactual in 2019, over 500 posterior draws, and the bar is its 90% credible band.

As Table 8 and Figure 48 show, the idiosyncratic labor-productivity shock has the largest median contribution to inequality. Removing it lowers the top-10% income share by 9.6 percentage points in China and 10.5 in Korea, and raises the bottom half's share by 5.1 and 4.0. The aggregate shock, by contrast, has contributions close to zero across outcomes, consistent with a shock common to all agents largely moving the distribution together.

Between the two preference dispersions, risk aversion has the larger median contribution. It contributes 3.1 points of China's top-10% income share and 3.5 of its wealth share, against 0.1 and 1.6 for the discount factor. The division suggests that risk aversion disperses labor supply, and so reaches income, in a way time preference does not.

The complementarity between capital and labor in the CES production function has a negative median contribution to inequality. With an elasticity below one, this result reflects the capital deepening of the catch-up period pushing income toward labor, the factor that grew relatively scarce, rather than toward capital. Had China's technology been Cobb-Douglas rather than its posterior median of 0.73, its top-10% income share would be 4.7 points higher, with 85% agreement.

Both taxes have small median contributions and no clearly determined effect on wealth concentration. Removing the capital tax raises China's top-10% wealth share by 0.4 points and Korea's by 0.3, and removing Korea's inheritance tax raises its own by 0.7. The limited response reflects the fact that the revenue funds state investment and public spending through

(22), with no transfer to households beyond the consumption floor.

5 Conclusion

This paper has studied growth and (re)distribution in China (1982–2019) and Korea (1962–2019) through the lens of a non-stationary heterogeneous-agent model formulated as a Bayes net. The CE-anchored quasi-rational equilibrium framework, and the computational pipeline built on amortized optimization and powered by deep learning, are reusable for the broader questions of growth and inequality in development economics.

The model collects engines that the growth-accounting and developmental-state literatures usually study apart. In particular, it explores the roles of technology catch-up, human capital accumulation, state-owned capital, and the social safety net in growth and distribution, which are rarely studied together in a heterogeneous-agent growth model. Agents differ in preferences as well as in idiosyncratic labor productivity and in what they inherit. To handle the computational complexity that comes with these rich features, I develop the CE-anchored quasi-rational equilibrium: agents forecast prices in an imagined deterministic auxiliary economy, attach subjective risk to that anchor, and re-forecast as the economy moves.

Bayesian estimation is feasible because every expensive solve is amortized into a neural network: decision rules cloned from an exact solver, anchor prices learned without labels from the market-clearing loss, and the posterior obtained by simulation-based Bayesian inference. On a single laptop, generating the training data takes less than 20 hours and training the networks takes less than 3 hours, with the resulting parameter uncertainty carried into every counterfactual experiment.

The counterfactuals traced out the role of each major channel in growth and distribution. The numerical effects reported here are posterior medians. Technology catch-up is first order for output, worth a 217% increase in China’s per-capita output and 306% in Korea’s by 2019, yet it has no clear effect on wealth shares, while the education expansion adds 3.6% and 16.6% to output by 2019 and 25.7% and 39.9% by 2040. The fertility decline has positive effects on per-capita output of 5.7% in China and 9.8% in Korea, while its distributional effects largely reflect composition: the missing cohorts would have been concentrated near the bottom as students and as young workers with the least accumulated human capital. Handing households a fifth of state investment reduces the top-10% wealth share by over 23 percentage points in both countries, at a cost of 22% of China’s output and 26% of Korea’s. Halving fiscal efficiency, which never reaches a household, has a larger output cost but small effects on inequality, with

no sign clearly determined. Raising the consumption floor to a fifth of the wage has output costs of 3.2% in China and 7.6% in Korea and consumption gains at the first percentile of 36% and 97%, although the corresponding credible intervals are wide. Among the estimated primitives, idiosyncratic labor-productivity risk has the largest contribution, worth 6.2 points of China’s top-10% wealth share and 8.6 of Korea’s, while the contributions of preference heterogeneity are smaller: risk-aversion dispersion contributes 3.5 points in each country and the discount factor 1.6 and 1.5.

The model is silent on several channels of growth and distribution that a richer one would carry. The economy is closed and the trade and capital flows play important roles in both countries’ economic growth. Households hold a single asset, so the housing and land booms that weigh heavily on measured wealth are absent. Other potentially important channels are missing too: China’s rural-urban income gap (Yang, 1999) and its urban housing reform of 1998, the repeal of Korea’s wage-suppression policies in the 1980s, and the rise of non-regular employment after the 1997 Asian financial crisis (Hong et al., 2025). Adding an open-economy block and a second asset such as housing is the natural next step for this line of work.

Appendix A Building Dataset

A.1 Population Data

I construct the entrant-cohort sizes $N_{j,t}$ and the cohort lifespans $d_{j,t}$ that drive the model’s demography. The inputs are total population and crude birth rates from the World Bank, the base-year population censuses (China 1982, Korea 1962), and period life expectancy at birth from OUR WORLD IN DATA.²¹ The construction has three parts: entrant cohorts, lifespans, and normalization.

Entrant cohorts. The entrant cohort $N_{j,t}$ is the mass of agents who reach adjusted age $a = 1$ (calendar age 15) in year t . I build it from two sources. Where the World Bank data are available, I take $N_{j,t}$ to be the size of the cohort born fifteen years earlier, total population times the crude birth rate in year $t - 15$. For the early cohorts whose birth year precedes the birth-rate series, I read $N_{j,t}$ from the base-year census, distributing each five-year age group uniformly across single ages and mapping a person of calendar age g in the census year to entry year $t_{s,j} - g + 15$. I rescale the birth-based series by a constant factor so that it joins the census

²¹<https://ourworldindata.org/grapher/life-expectancy>.

counts continuously at the splice, matching them over the first eight years in which the two overlap (the entire overlap when it is shorter).

Cohort lifespans. A cohort entering at t remains in the model for $d_{j,t}$ adjusted-age years and then exits. I set $d_{j,t} = e_{j,t}^0 - 14$, rounded to the nearest year, where $e_{j,t}^0$ is period life expectancy at birth. I fill the sparse early years of the life-expectancy series by linear interpolation. Because life expectancy falls in some years, gaps in the age pyramid are permitted.

Normalization. The model’s aggregates are homogeneous of degree one in the cohort path $\{N_{j,t}\}$, so their overall scale is not identified. I divide the entire path by the base-year population—the cohorts alive at the catch-up start $t_{s,j}$ ($t_{s,cn} = 1982$, $t_{s,kr} = 1962$)—so that base-year population equals one. The resulting panel records the mass of every alive cohort by adjusted age in each year.

A.2 Education Data

I assign each entrant cohort $N_{j,t}$ highest-attainment shares over the three education levels of [Section 2.1](#), below high school (e_b), high school (e_h), and college (e_c), which attach directly to the normalized population panel of [Section A.1](#). Cohorts whose graduation flows are observed (China from the 1967 entry cohort, Korea from 1968) take their shares from those flows; earlier cohorts take a backcast path fitted to the BARRO-LEE attainment series.

Data sources. Graduation counts come from CEIC, except Chinese senior high-school graduates (National Bureau of Statistics) and Korean junior-college graduates (Korean Educational Development Institute). High-school graduates are the senior plus vocational counts for China and the general plus vocational counts for Korea; college graduates are the regular undergraduate plus junior-college counts for China and the four-year university plus junior-college counts for Korea. Counting junior colleges as college on both sides keeps the category all-tertiary, so it is comparable across countries and with the BARRO-LEE tertiary-completed shares that anchor the pre-flow backcast, which CEIC data doesn’t cover. The educational attainment of the population aged fifteen and over comes from the BARRO-LEE dataset ([Barro and Lee, 2013](#)), reported at five-year intervals.

Flow cohorts. The flow cohorts are those the graduation counts reach, mainly the CEIC series. A cohort entering at t completes high school at $t + 3$ and college at $t + 7$, so its college share is the number of college graduates in year $t + 7$ divided by the entrant cohort $N_{j,t}$, and its

high-school share is the analogous ratio net of those who continue to college, so that each level counts the *highest* attained. The uneducated share is the residual, clipped at zero, and beyond the end of the graduation series I extrapolate using the last observed value.

Pre-flow cohorts. For the cohorts entering before the flows begin I take the composition from the BARRO-LEE attainment of the fifteen-and-over population (Barro and Lee, 2013). Its categories are not level-identical to the graduation-flow concept, so I calibrate it against the flow data, estimating a definitional scale κ per series together with a backcast path that splices into the first flow cohort without a seam. The Korean scales are close to one; for China, BARRO-LEE completed secondary is roughly twice the vocational-plus-senior graduation concept ($\kappa_h \approx 0.47$).

A.3 Macroeconomic Time Series Data

I build the model’s macroeconomic inputs—a common technology frontier and each country’s output and capital—from the PENN WORLD TABLE 11.0 (Feenstra et al., 2015), all at constant national prices.

Technology frontier. I use US total factor productivity (RTFP^{NA}) to discipline the world technology frontier $A_{f,t}$ of (2): its average growth informs the prior on the frontier growth rate g_f , while matching 1962 Korean output centers the prior on the level $A_{f,1962}$ (Section B). The frontier path itself is the estimated constant-growth process, not the raw index.

Output and capital. Real GDP (RGDP^{NA}) maps to output $Y_{j,t}$ in the production function (1), and the capital stock (RN^{NA}) to the economy’s total physical capital—the additive sum $K_{p,j,t} + K_{s,j,t}$ of private and state capital. The state-capital share divides this total into its two components, which the production function bundles into the effective aggregate $K_{j,t}$ through (7). To keep these aggregates consistent with the population panel I express both series in trillions of 2021 USD and divide them by the base-year population in billions of people—the catch-up-start cohorts that Section A.1 normalizes to one ($t_{s,cn} = 1982$, $t_{s,kr} = 1962$). Output and capital are homogeneous of degree one in the population, so dividing the whole path by this single constant leaves the capital-output ratio unchanged and expresses each series in thousands of 2021 USD per base-year person.

A.4 Inequality Data

For each country I take, from the WORLD INEQUALITY DATABASE (Chancel et al., 2022), the equal-split adult shares of pre-tax national income and of net personal wealth held by the top 1%, the top 10%, and the bottom 50% of adults. Pre-tax income is labor plus capital income before direct taxes and transfers (with pensions counted on a distribution basis), and wealth is household net worth: all non-financial and financial assets minus debts. Net worth is the closest available counterpart to the agents' predetermined saving $k_{i,a}$, the stock (39) accumulates out of income not consumed. Both are measured across adults with each couple's resources split equally between partners (WID's equal-split series), which is the closer match to the model's individual agents. The shares are unit-free and enter the estimation as targets, so they need no normalization.

China's income and wealth shares (1982–2019) and Korea's income shares (1962–2019) are fully observed at annual frequency. Korean wealth shares are reported only for 1960, 1970, and 1980–2019; I fill the missing pre-1980 years (1962–1969 and 1971–1979) by linear interpolation between the nearest observed years.

A.5 Labor Income Tax Rates Data

I assemble the progressive labor-income tax schedule $\mathcal{T}_{i,j,t}$ of (18)—the marginal rates and bracket thresholds in force each year—for each country.

China's schedule follows the Individual Income Tax Law of the People's Republic of China, from its 1980 enactment through the 1993 and 2011 revisions, with threshold-only amendments in 2007 and 2018: the top marginal rate stays at 45% throughout, while the number of brackets moves from seven (1980–1992) to ten (1993–2010) to eight (2011–2019), and the 2011 reform cuts the lowest non-zero rate from 5% to 3%. Korean rates for 1960–1974 are proxied by the 1974 schedule (Yoo, 2000) and for 1975–2019 come from `taxwatch.co.kr`; over successive revisions the top rate falls from 70% (1960–1981) to 35% (2005–2011) before rebounding to 42% (2018–2019), the bracket count compresses from sixteen to as few as four, and the lowest non-zero rate stays within a 5–10% band, settling at 6% by 2019.

I convert each bracket threshold from local currency to thousands of 2021 USD—the model's individual-monetary unit—using annual exchange rates from the Federal Reserve (with CEIC covering the years before 1981) and the US consumer price index, also from the Federal Reserve.

This exact schedule is the labor income tax liability $\mathcal{T}_{i,j,t}$ throughout: the tax owed on gross labor income, with net income gross minus this liability. Appendix C describes how it is

summarized into features for the surrogate networks and perturbed across training draws.

Appendix B Prior Distributions

This appendix gives the rationale for the prior distributions summarized in [Table 1](#). Throughout, $\mathcal{TN}(\mu, \sigma^2; L, U)$ denotes a normal distribution with mean μ and standard deviation σ truncated to $[L, U]$, and $\text{TPois}(\lambda; L, U)$ a Poisson distribution with rate λ truncated to the integers $\{L, \dots, U\}$.

B.1 Frontier and TFP Adoption

This group parameterizes the world technology frontier and each country’s catch-up toward it, as specified in (2)–(6).

- **Frontier TFP growth rate g_f**

$$g_f \sim \mathcal{TN}(0.0055, 0.0015^2; 0.002, 0.01).$$

I center the prior at the average annual growth of US TFP over the 1962–2019 sample, about 0.55%. The 5th to 95th percentiles of the truncated distribution, about $[0.32\%, 0.80\%]$ —corresponds to constant-growth frontier paths whose implied 2019 level, compounded from the 1962 base, lies roughly 12% below to 15% above the observed US TFP level. The prior is therefore consistent with the realized frontier while leaving room for forecast uncertainty. I truncate the support to $[0.2\%, 1\%]$: the lower bound rules out a near-stagnant frontier, which is inconsistent with the sustained TFP growth in the data, and the upper bound rules out implausibly fast frontier growth and keeps the forward simulation numerically stable.

- **Frontier TFP level $A_{f,1962}$**

$$A_{f,1962} \sim \mathcal{TN}(28, 8.3^2; 3, 53).$$

The level of the frontier sets the monetary units of output, which I do not normalize away: because the labor-income tax is specified on income in thousands of 2021 USD, model income must be in the same units for the schedule to apply directly. I therefore pin the level by matching 1962 Korean output. At the catch-up start the country TFP is $A_{kr,1962} = \theta_{A,kr} \lambda_{kr,1962} A_{f,1962}$, and output (1) is $Y = A_{kr,1962} F(K, L)$ with F the CES

capital–labor bundle, so

$$A_{f,1962} = \frac{Y_{kr,1962}}{\theta_{A,kr} \lambda_{kr,1962} F(K_{kr,1962}, L_{kr,1962})},$$

evaluated at the data capital and output, the prior means of θ_A , λ , and the production parameters, and the aggregate effective labor L (with mean hours near one third and $h \approx 1$). This centers the prior near 28; the standard deviation is generous and the support wide to span the uncertainty in θ_A , λ , and the effective-labor aggregate.

- **Technology ceilings** $\theta_{A,cn}$, $\theta_{A,kr}$

$$\theta_{A,cn} \sim \mathcal{TN}(0.8, 0.15^2; 0.5, 1),$$

$$\theta_{A,kr} \sim \mathcal{TN}(0.8, 0.15^2; 0.7, 1)$$

In (3) the ceiling $\theta_{A,j}$ is the asymptotic fraction of the frontier that country j can reach, reflecting institutional and structural limits. I center both priors at 0.8, encoding the belief that each country can substantially close its gap to the frontier, with a large standard deviation of 0.15 for the deep uncertainty about this long-run outcome. The lower bounds differ by country, set at or near each country’s highest attained TFP-to-frontier ratio in PWT 11.0: China has reached about 47% of the US level by 2023, giving a floor of 0.5, while Korea has reached about two-thirds, peaking at 71% in 2002, giving a higher floor of 0.7. Each floor holds $\theta_{A,j}$ essentially at or above the ratio the country has already attained, so that $\lambda_{j,t}$ in (4) stays at or below one; at Korea’s floor, which rounds its 2002 peak down, λ touches one in that single year.

- **Initial TFP-to-ceiling ratios** $\lambda_{cn,1982}$, $\lambda_{kr,1962}$

$$\lambda_{cn,1982} \sim \mathcal{TN}(0.2, 0.05^2; 0.1, 0.5),$$

$$\lambda_{kr,1962} \sim \mathcal{TN}(0.3, 0.05^2; 0.1, 0.5)$$

The priors for the initial TFP positions relative to the country-specific ceiling, $\lambda_{j,t_{s,j}}$ in (3), are set from the historical context of each country at its catch-up start year $t_{s,j}$.

For China, the catch-up begins in $t_{s,cn} = 1982$. By this time, China had a foundational industrial base from the pre-reform era (Walder, 2015, Chapter 14) and had begun its agricultural reforms (Lin, 1992). This suggests a starting TFP that was neither negligible nor advanced. The PWT 11.0 CTFP (TFP at current purchasing power parities relative to the US) places China at 19.1% of the US level in 1982 (assuming $\theta_{A,cn} = 1$), so I center

the prior at 0.2.²²

For Korea, the catch-up starts in $t_{s,kr} = 1962$. The country was still recovering from the havoc of the Korean War, with industrial production having only recently returned to pre-war levels (Kohli, 2004). The PWT 11.0 CTFP of 26.7% in 1962 (assuming $\theta_{A,kr} = 1$) suggests Korea started at a somewhat higher TFP share of its ceiling than China, so I center the prior at 0.3. The observed CTFP equals $\theta_{A,j} \lambda_{j,t_{s,j}}$, so reading each center off it takes $\theta_{A,j} = 1$, while the ceiling prior is centered at 0.8. The standard deviation of 0.05 is set to absorb this, so the prior spans the joint uncertainty in $\lambda_{j,t_{s,j}}$ and $\theta_{A,j}$ rather than committing to either anchor.

The lower bounds are held strictly above zero because the catch-up law (4) operates on $\log \lambda_{j,t}$.

- **Maximum adoption speed** v_{max}

$$v_{max} \sim \mathcal{TN}(0.12, 0.05^2; 0.05, 0.25)$$

This is the adoption speed of a fully college-educated, full-time workforce, reached as the education-weighted density $\mathcal{H}_{j,t}$ of (5) approaches one. Because the density is well below one along the transition (education weights concentrate mass in the small college group), the realized speed stays far below v_{max} . At the center of 0.12, a workforce held at $\mathcal{H}_{j,t} = 1$ would take about 27 years to carry China from $\lambda_{cn,1982} = 0.2$ to $\lambda = 0.95$, and about 25 years to carry Korea from 0.3, roughly one working generation. The truncation at 0.25 rules out closing the gap in ten years, and the floor of 0.05 stretches the same transition beyond sixty years.

- **High-school adoption weight** ω_h

$$\omega_h \sim \mathcal{TN}(1.5, 0.5^2; 1, 4)$$

This is the weight of high-school workers' human capital in the education-weighted density of (5), relative to the below-high-school weight normalized to one. The lower bound of one imposes that schooling never lowers a worker's contribution to absorptive capacity; the center of 1.5 is a mild premium consistent with secondary schooling aiding imitation.

²²PWT 10.01 reported a much higher figure of 37.4% and, anomalously, implied China's average TFP share was higher during 1957–1976 (47.8%) than during the reform era of 1977–2019 (39.6%). PWT 11.0 resolves this inconsistency: the pre-reform average (22.1%) is now well below the reform-era average (31.7%), consistent with the literature that TFP growth is a major driver of China's post-reform economic success (Zhu, 2012).

- **College adoption-weight increment $\Delta\omega_c$**

$$\Delta\omega_c \sim \mathcal{TN}(2.0, 1.5^2; 0, 6)$$

Rather than estimate the college weight directly, I estimate its non-negative increment over the high-school weight, so that $\omega_c = \omega_h + \Delta\omega_c$ and the ordering $\omega_c \geq \omega_h \geq 1$ holds by construction. Against the center of 1.5 for ω_h , the center of 2.0 puts the college weight near 3.5: where secondary schooling aids imitation, tertiary schooling is what carries innovation near the frontier (Vandenbussche et al., 2006). The standard deviation of 1.5 and the support up to 6 hold that belief weakly, and the floor at zero leaves room for no college premium at all.

- **Country TFP-shock persistence $\eta_{j,u}$**

$$\eta_{j,u} \sim \text{Beta}(4, 2)$$

This is the persistence of the AR(1) TFP shock in (6). The shock should be persistent (the prior mean is two-thirds) but mean-reverting rather than a unit-root process, so I use a Beta(4, 2) prior, which discourages values too close to 0 or 1.

- **Country TFP-shock std. dev. $\sigma_{j,u}$**

$$\sigma_{j,u} \sim \mathcal{TN}(0.05, 0.025^2; 0, 0.1)$$

This is the standard deviation of the innovations to the TFP shock in (6). The relevant consideration is the scale of the observation equation: a shock of 0.1 to log TFP would be a large perturbation to annual log output, so I constrain the prior to the small range $[0, 0.1]$, with the nonnegative lower bound enforcing a valid standard deviation.

B.2 Nested CES Production

The production technology has two CES nests: an outer final-output nest that combines aggregate capital and efficient labor in (1), and an inner aggregator that combines private and state capital in (7). The priors below reflect this structure together with theoretical restrictions and numerical safety.

- **Final-output CES share** $\alpha_{Y,j}$

$$\alpha_{Y,j} \sim \mathcal{TN}(0.5, 0.15^2; 0.05, 0.95)$$

$\alpha_{Y,j}$ is the CES weight on capital in (1). I give it a symmetric, weakly informative prior centered at 0.5; the truncation away from the endpoints keeps both factors present and the marginal products in (11) well behaved.

- **Capital-labor elasticity** $\sigma_{Y,j}$

$$\log \sigma_{Y,j} \sim \mathcal{TN}(\log 0.6, 0.25^2; \log 0.3, \log 1.3)$$

I place a truncated-normal prior on $\log \sigma_{Y,j}$, so that the prior is symmetric in ratios rather than in levels. I center it at $\sigma_{Y,j} = 0.6$, below the Cobb-Douglas value of one: the weight of the evidence is that capital and labor are gross complements. Chirinko (2008) reads the estimates as pointing to 0.4–0.6, and the meta-analysis of Gechert et al. (2022), covering 121 studies, puts the raw literature mean at 0.3 once the publication bias is addressed. A prominent minority instead finds $\sigma > 1$: Karabarbounis and Neiman (2014) estimate 1.25, reading the global labor-share decline as capital substituting for labor, so I keep a thin tail above one rather than imposing complementarity. Working in logs is what delivers that tail: log-symmetry makes a given factor above the center as likely as the same factor below, placing about 2% of the prior mass above one. I cap the support at 1.3 because for $\sigma_{Y,j} > 1$ to avoid the nest becomes asymptotically linear in capital. The lower bound of 0.3 excludes a near-Leontief technology and coincides with the fully corrected meta-analytic mean.

- **Private-capital share** $\alpha_{K,j}$

$$\alpha_{K,j} \sim \mathcal{TN}(0.5, 0.2^2; 0.1, 0.95)$$

$\alpha_{K,j}$ is the CES weight on private capital in the aggregator (7). I give it a symmetric, weakly informative prior centered at 0.5; the truncation away from the endpoints keeps both private and state capital present and their shares in (10) interior.

- **Capital aggregator elasticity** $\sigma_{K,j}$

$$\sigma_{K,j} \sim \mathcal{TN}(0.5, 0.2^2; 0.1, 0.95)$$

The elasticity between private and state capital is restricted to $\sigma_{K,j} < 1$ so that the two are gross complements, consistent with the model’s premise that infrastructure and other state capital raise the productivity of private investment while having limited value on their own. I center the prior at one half, comfortably inside the complementarity region, with a spread that allows substantial uncertainty while the upper truncation at 0.95 keeps the aggregator safely inside the gross-complement region $\sigma_{K,j} < 1$ and the lower bound avoids a near-Leontief aggregator.

- **Depreciation rate δ**

$$\delta \sim \mathcal{TN}(0.05, 0.01^2; 0.01, 0.1)$$

δ is the common annual depreciation rate of private and state capital in (11). I center it at five percent, close to the 4.8% annual rate (a quarterly 1.2%) that Cooley and Prescott (1995) calibrate for a broad capital aggregate in the real-business-cycle tradition, which suits the model’s broad private-and-state capital. The spread is kept tight—and capped at ten percent—to keep the net capital returns $r_{p,j,t}$ and $r_{s,j,t}$ in (11) positive.

B.3 Government Policy

These parameters govern the fiscal block in (22), fixing how generously the state insures consumption, how heavily it taxes capital and bequests, how efficiently it converts revenue into spending, and how costly public education is.

- **China consumption-floor ratios $\Lambda_{cn,1}, \Lambda_{cn,2}$**

$$\Lambda_{cn,1} \sim \mathcal{TN}(0.03, 0.02^2; 0, 0.2),$$

$$\Lambda_{cn,2} \sim \mathcal{TN}(0.1, 0.04^2; 0.02, 0.3)$$

These set the consumption floor in (14) as a ratio to the current wage $w_{j,t}$ before and after the structural break in welfare policy at $t_{g,cn} = 2012$; the priors are informed by historical data on China’s social safety net. Because $w_{j,t}$ is a full-time worker’s labor income while observed average income reflects mean hours near one third, a floor quoted against average income maps to roughly one third as large a ratio to $w_{j,t}$. Before 1993, urban social security was largely tied to employment in state-owned enterprises, while rural coverage was sparse (Hussain, 1994). The government began establishing a minimum living security system (*Dibao*) in urban areas in 1993, extending it to rural areas only after 2007, and benefit levels were initially low: reports from the Ministry of Civil Affairs indicate the

rural benefit was just 58 yuan per month in 2003, and by 2011 it stood at about 56% of average rural labor income and 19% of average urban labor income. Because coverage was patchy and the benefit low in absolute terms, I center the pre-2012 prior at about 0.03 of $w_{j,t}$, reflecting a thin early safety net rather than a direct mapping of these ratios. In the post-2012 period, welfare support expanded significantly, notably with Targeted Poverty Alleviation: by 2020 the rural minimum living standard (Dibao) reached 5,482 yuan, about 32% of rural per-capita disposable income (17,131 yuan).²³ This is about 0.1 of the full-time wage $w_{j,t}$, so I raise the prior center to 0.1 and widen the support to allow for a much more aggressive post-reform regime.

- **Korea consumption-floor ratios** $\Lambda_{kr,1}$, $\Lambda_{kr,2}$

$$\Lambda_{kr,1} \sim \mathcal{TN}(0.03, 0.02^2; 0, 0.2),$$

$$\Lambda_{kr,2} \sim \mathcal{TN}(0.1, 0.04^2; 0.02, 0.3)$$

For Korea the structural break is set at $t_{g,kr} = 1988$. Before 1988, during its high-growth period, Korea's social welfare expenditure was relatively low, so I assign a prior similar to China's early period. After 1988, Korea established a more universal welfare system (Koh, 2015). The livelihood benefit under the Basic Livelihood Security Programme carries an income criterion of 30% of median income (Hwang et al., 2022). Applying the same full-time-wage conversion as for China, this is roughly 0.1 of $w_{j,t}$, so I center the post-break prior at 0.1, in line with China's post-2012 floor, while keeping the same wide support to allow for a more generous regime.

- **Capital income tax rates** $\tau_{k,cn}$, $\tau_{k,kr,1}$, $\tau_{k,kr,2}$

$$\tau_{k,cn} \sim \mathcal{TN}(0.10, 0.05^2; 0, 0.3),$$

$$\tau_{k,kr,1} \sim \mathcal{TN}(0.10, 0.05^2; 0, 0.3),$$

$$\tau_{k,kr,2} \sim \mathcal{TN}(0.20, 0.05^2; 0, 0.3)$$

As discussed in Section 2.3, the effective capital-income tax rate in (17) is a single constant for China, $\tau_{k,cn,t} \equiv \tau_{k,cn}$, and piecewise constant for Korea, breaking at the 1997 Asian financial crisis into a pre-crisis rate $\tau_{k,kr,1}$ ($t \leq 1997$) and a post-crisis rate $\tau_{k,kr,2}$ ($t > 1997$). In the global effective-tax-rate database of Bachas et al. (2026), effective capital taxation is markedly lower in emerging economies than in high-income ones, though

²³Rural subsistence-allowance standard (People's Daily): <http://politics.people.com.cn/n1/2020/1229/c1001-31981848.html>; rural per-capita disposable income (National Bureau of Statistics): https://www.stats.gov.cn/sj/zxfb/202302/t20230203_1900970.html.

it has been rising: across developing countries the effective rate rose from 10% in 1989 to 18% in 2018, and China’s from 6% in 1994, when it enters their sample, to 24% in 2018, while in high-income countries it declined from 38–39% in the late 1960s to 32–33% in the late 2010s. I therefore center China’s rate and Korea’s pre-crisis rate at 10%, in the lower part of the range these emerging-economy series span, and Korea’s post-crisis rate at 20%, reflecting the higher effective capital taxation of its post-crisis tax regime, with a common standard deviation of 0.05 and a shared upper bound of 30%.

- **Korea’s transfer tax rate** $\tau_{m,kr}$

$$\tau_{m,kr} \sim \mathcal{TN}(0.3, 0.1^2; 0.1, 0.5)$$

As specified in (21), China has no inheritance tax ($\tau_{m,cn,t} = 0$), while Korea’s transfer tax rate is constant over time. Korea’s statutory inheritance tax rate is among the highest in the OECD, with a top marginal rate of 50%. I therefore center the prior at 0.3 to reflect a high but not maximal effective rate on the bequests taxed in (20), with bounds capturing the plausible range.

- **Fiscal efficiency** $\zeta_{j,1}, \zeta_{j,2}$

$$\zeta_{j,1} \sim \mathcal{TN}(0.7, 0.1^2; 0.5, 1),$$

$$\zeta_{j,2} \sim \mathcal{TN}(0.8, 0.1^2; 0.5, 1)$$

The efficiency parameter in (23) captures the fraction of total government revenue that is effectively deployed for public spending in the budget constraint (22). It is modeled as piecewise constant with a structural break at $t_{g,j}$, corresponding to Korea’s democratic transition and China’s anti-corruption campaign. For the pre-reform period, I expect moderate but imperfect efficiency due to less mature institutions and weaker oversight; for the post-reform period, I expect a further improvement. Both priors restrict the support to $[0.5, 1]$ to rule out implausibly low deployment rates that would drive the posterior toward degenerate solutions under sparse summary statistics.

- **Students-to-teachers ratio** θ_κ

$$\theta_\kappa \sim \mathcal{TN}(20, 5^2; 10, 30)$$

The parameter θ_κ in (16) represents the number of students per teacher, which determines the per-student education cost $w_{j,t}/\theta_\kappa$. Because (16) charges only the high-school and

college cohorts, the relevant ratios are those of senior secondary and tertiary education. For 2024 the Chinese Ministry of Education reports 12.7 students per teacher in regular senior high schools, 16.9 in secondary vocational schools, 17.1 in regular undergraduate institutions, and 18.5 in higher vocational colleges.²⁴ I center the prior at 20, just above this range to allow for the larger cohorts of earlier decades, truncating to (10, 30) to exclude both implausibly small classes and extreme overcrowding.

B.4 Preferences and Family Heterogeneity

Each preference parameter is family-specific, drawn from a population distribution whose mean and log standard deviation carry their own hyperpriors; the items below justify the center, spread, and truncation of each draw in light of the period utility (24) and the bequest term.

- **Discount factor** β_i

$$\begin{aligned}\beta_i &\sim \mathcal{TN}(\mu_\beta, \sigma_\beta^2; 0.5, 1), \\ \mu_\beta &\sim \mathcal{TN}(0.95, 0.05^2; 0.8, 1), \\ \ln \sigma_\beta &\sim \mathcal{N}(-2.5, 0.5^2)\end{aligned}$$

The family discount factor is truncated to $[0.5, 1]$ to enforce impatience over a finite life cycle while ruling out economically implausible extremes. The population mean is centered at 0.95, consistent with conventional annual patience, and bounded within $[0.8, 1]$. Because lifetime wealth is acutely sensitive to β_i near one, the hyperprior on $\ln \sigma_\beta$ keeps the cross-family spread weakly informative rather than presupposing it: it centers the realized standard deviation at about 0.06 for β_i , with most families between 0.80 and 0.99—while its tail reaches down to the truncation floor.

- **Risk aversion** γ_i

$$\begin{aligned}\gamma_i &\sim \mathcal{TN}(\mu_\gamma, \sigma_\gamma^2; 0.5, 5), \\ \mu_\gamma &\sim \mathcal{TN}(2.5, 0.5^2; 1.5, 3.5), \\ \ln \sigma_\gamma &\sim \mathcal{N}(0, 0.2^2)\end{aligned}$$

The coefficient of relative risk aversion in (24) is truncated to $[0.5, 5]$: this spans the range commonly used in life-cycle work and admits a mild sub-logarithmic, higher-intertemporal-substitution region $\gamma_i \in [0.5, 1)$, while excluding the near-risk-neutral region $\gamma_i < 0.5$, whose discounted value explodes under high-return price paths. The family mean μ_γ is centered at 2.5 (moderate risk aversion) and itself truncated to $[1.5, 3.5]$. The hyperprior

²⁴ *Statistical Bulletin on National Education Development in 2024*, http://www.moe.gov.cn/jyb_sjzl/sjzl_fztjgb/202506/t20250611_1193760.html.

on $\ln \sigma_\gamma$ is centered at zero, so the implied cross-family standard deviation centers on 1, allowing moderate heterogeneity in risk attitudes across families; its tight spread keeps that dispersion from collapsing to homogeneity or exploding to extreme curvature draws.

- **Bequest weight** ψ_i

$$\begin{aligned}\psi_i &\sim \mathcal{TN}(\mu_\psi, \sigma_\psi^2; 0, 2), \\ \mu_\psi &\sim \mathcal{TN}(1, 0.25^2; 0.5, 2), \\ \ln \sigma_\psi &\sim \mathcal{N}(-0.7, 0.3^2)\end{aligned}$$

The bequest weight in the terminal utility term is truncated to $[0, 2]$, with the lower bound permitting the no-bequest case. The mean is centered near unity so that the bequest motive is comparable in strength to annual consumption utility. The hyperprior on $\ln \sigma_\psi$ centers the cross-family dispersion at $\exp(-0.7) \approx 0.5$ —half the mean weight—so that weak-bequest families are genuinely represented: ψ_i near zero lies only two standard deviations below the mean, and about one family in seven draws $\psi_i < 0.5$, which allows a sizable fraction of households that the empirical bequest literature finds saving with little bequest motive (De Nardi et al., 2010).

- **Labor-disutility weight** $\log \chi_i$

$$\begin{aligned}\log \chi_i &\sim \mathcal{N}(\mu_{\log \chi}, \sigma_{\log \chi}^2), \\ \mu_{\log \chi} &\sim \mathcal{N}(-1.92, 1^2), \\ \ln \sigma_{\log \chi} &\sim \mathcal{N}(-0.7, 0.5^2)\end{aligned}$$

The weight χ_i enters (24) as a pure *level* parameter. I calibrate it on the agent’s dynamic problem: I solve the certainty-equivalent problem for a representative agent, with every other parameter at its prior mean, and recenter $\mu_{\log \chi}$ until optimal working time near calendar age thirty sits at about one third of the time endowment, interior to $(0, 1)$. This yields $\mu_{\log \chi} \approx -1.92$, a median weight $\chi \approx 0.15$. The wide prior standard deviation reflects the uncertainty in this calibration procedure, and the log standard deviation lets families differ in their distaste for work.

- **Labor-disutility curvature** ν_i

$$\begin{aligned}\nu_i &\sim \mathcal{TN}(\mu_\nu, \sigma_\nu^2; 0.2, 4), \\ \mu_\nu &\sim \mathcal{TN}(1.0, 0.5^2; 0.2, 4), \\ \ln \sigma_\nu &\sim \mathcal{N}(-1.0, 0.4^2)\end{aligned}$$

The convexity ν_i sets the Frisch elasticity of labor supply, $1/\nu_i$ —the response of hours to the wage holding the marginal utility of wealth fixed—and is truncated to $[0.2, 4]$, a Frisch elasticity in $[0.25, 5]$, excluding both near-linear and extremely steep disutility. Centering the mean at 1.0 with a modest hierarchical spread gives a central 90% prior interval $\nu_i \in [0.34, 2.1]$, a Frisch elasticity of roughly $[0.5, 3]$.

B.5 Labor Productivity and Human Capital

The idiosyncratic labor-productivity shock in (27) and the human-capital law of motion in (31) share several parameters whose priors are justified here.

- **Persistence of labor shock η_ε**

$$\eta_\varepsilon \sim \text{Beta}(18, 2)$$

This is the persistence of the AR(1) shock $\varepsilon_{\iota,a}$ in (27). As with the other persistence parameters in Table 1, I use a Beta prior; here the shape is $(18, 2)$, whose mean is 0.9 and whose central 90% interval is $[0.77, 0.98]$. This spans the range the income-process literature estimates for the persistent component of idiosyncratic earnings, from 0.91 in Flodén and Lindé (2001) to 0.97 in Krueger et al. (2016).

- **Std. dev. of labor shock σ_ε**

$$\sigma_\varepsilon \sim \mathcal{TN}(0.18, 0.06^2; 0.05, 0.4)$$

This scales the multiplicative innovation $\varepsilon_{\varepsilon,\iota,a}$, which enters efficiency units through $\exp(\varepsilon_{\iota,a})$ in (26). The innovation mean is not a free parameter—it is pinned to $-\sigma_\varepsilon^2/(2(1 + \eta_\varepsilon))$ by the unit-mean restriction in (28), so this prior governs the dispersion only. The center of 0.18 is near the estimated standard deviation of the persistent earnings innovation, 0.21 in Flodén and Lindé (2001) and 0.20 in Krueger et al. (2016).

- **High-school premium Δh_h**

$$\Delta h_h \sim \mathcal{TN}(0.45, 0.12^2; 0.21, 0.75)$$

With $h_{e_b} = 1$ fixed by normalization, high-school entry human capital is $h_{e_h} = 1 + \Delta h_h$ in (29), so this is the high-school premium in starting human capital. I anchor it to the model’s own learning technology: under (31) at the learning prior means, an uneducated

agent who worked full time through the three years of high school would gain only about 0.12 of human capital, so any premium below that would leave high school not worth its forgone experience. The lower bound of 0.21—the gain from five such years of work—keeps high school strictly profitable, and the mean of 0.45 centers the premium near ten years of work-equivalent experience.

- **College premium Δh_c**

$$\Delta h_c \sim \mathcal{TN}(0.55, 0.15^2; 0.24, 0.9)$$

College entry is built as $h_{e_c} = h_{e_h} + \Delta h_c$ in (29), so estimating this positive increment enforces the ordering $h_{e_b} < h_{e_h} < h_{e_c}$ by construction. The work-equivalent floor is set so that the degree pays off twice over: with both premia at their lower bounds, college entry human capital is 1.45, above both the 1.30 an uneducated agent would reach by working the seven years a college entrant forgoes and the 1.41 a high-school graduate would reach by working the four years between high-school and college entry. The mean of 0.55 places h_{e_c} near 2.0, the human capital an uneducated agent would reach after about twenty years of full-time work, so schooling years build the stock substantially faster than early working years.

- **Human-capital depreciation δ_h**

$$\delta_h \sim \mathcal{TN}(0.02, 0.015^2; 0, 0.08)$$

This is the per-period depreciation of human capital in (31). The small mean of 0.02 captures gradual skill obsolescence over a working life, and the $[0, 0.08]$ truncation keeps depreciation nonnegative while bounding it well below the learning-by-doing gains so that human capital does not erode implausibly fast.

- **Learning-by-doing scale $\gamma_{h,0}$**

$$\gamma_{h,0} \sim \mathcal{TN}(0.06, 0.03^2; 0, 0.2)$$

This scales the learning-by-doing term in (31). The center of 0.06 permits substantial early-career accumulation, and the upper bound of 0.2 prevents explosive human-capital growth that would dominate the net accumulation factor $(1 - \delta_h) + \gamma_{h,0}[\cdot]g_h(\ell)$.

- **Learning-age decline** $\gamma_{h,1}$

$$\gamma_{h,1} \sim \mathcal{TN}(0.008, 0.004^2; 0, 1/70)$$

This governs the linear decline in learning ability with age, $a - a_{e_t}$, the time since education completion, in (31). The very small mean of 0.008 makes learning ability fade gently. The upper bound $1/70$ is the reciprocal of the longest cohort lifespan in the sample, $\max_{j,t} d_{j,t_0} = 70$ adjusted-age years (calendar age 84): since the age penalty $a - a_{e_t}$ never exceeds the cohort's lifespan d_{j,t_0} over the working life, the bracketed term $1 - \gamma_{h,1}(a - a_{e_t})$ stays strictly positive, so that working never mechanically subtracts human capital.

- **Learning concavity** α_h

$$\alpha_h \sim \text{Beta}(2, 2)$$

This is the concavity of the normalized learning function $g_h(\ell)$. The $\text{Beta}(2, 2)$ prior is symmetric and confined to $(0, 1)$, which is the range that delivers diminishing returns to working time in human-capital accumulation while remaining agnostic about the precise degree of concavity.

- **Learning smoothing** $\mathfrak{d}_h$

$$\mathfrak{d}_h \sim \mathcal{TN}(0.1, 0.05^2; 0.01, 0.3)$$

This smoothing parameter shifts the argument of g_h so that the marginal effect of working time stays finite at $\ell_{\iota,a} = 0$, avoiding the unbounded slope that pure power functions exhibit at the origin. The mean of 0.1 provides mild regularization, with the strictly positive lower bound of 0.01 required for $\mathfrak{d}_h > 0$ and the upper bound of 0.3 keeping the learning function from becoming nearly linear.

B.6 Subjective Price Risk and CE Prediction

The subjective price risk that scales the CE anchors in (61), a bivariate AR(1) (62) with the correlated Gaussian innovations of (63), contributes the parameters whose priors are justified here.

- **Persistence of wage risk** $\eta_{w,j}$

$$\eta_{w,j} \sim \text{Beta}(4, 2)$$

The persistence of the log wage-risk component of $\xi_{j,t \rightarrow t'}$ in (62). The prior mass sits well above one-half so that a subjective wage surprise decays gradually rather than dissipating within a single year, consistent with the slow-moving nature of labor-market conditions, while the density still vanishes at unity to keep the AR(1) process stationary around the anchor.

- **Persistence of return risk $\eta_{r,j}$**

$$\eta_{r,j} \sim \text{Beta}(4, 2)$$

The persistence of the log return-risk component in (62). It shares the same weakly informative shape as $\eta_{w,j}$, encoding persistent but mean-reverting return surprises; the upper bound at one is informative here because a unit root would let subjective risk drift permanently away from the deterministic CE anchor.

- **Wage-risk innovation s.d. $\sigma_{w,j}$**

$$\sigma_{w,j} \sim \mathcal{TN}(0.05, 0.025^2; 0, 0.15)$$

The standard deviation of the log wage-risk innovation in (63). Since the process starts from $\xi_{j,t \rightarrow t} = \mathbf{0}$ and fans out with the forecast horizon, what matters is the unconditional standard deviation of the log price-risk process, $\sigma/\sqrt{1-\eta^2}$, which is proportional risk because the anchor is scaled by $\exp(\xi)$. Centering the innovation near 0.05 gives an unconditional standard deviation of about 0.07, a modest spread reflecting that the CE anchor already absorbs the predictable component of wages; the truncation keeps the standard deviation strictly positive and caps it at 0.15, near twenty percent unconditionally, which rules out implausibly volatile subjective wage paths and protects the covariance matrix from extreme draws.

- **Capital Return risk innovation s.d. $\sigma_{r,j}$**

$$\sigma_{r,j} \sim \mathcal{TN}(0.06, 0.03^2; 0, 0.18)$$

The standard deviation of the log return-risk innovation in (63). I center it just above $\sigma_{w,j}$, at 0.06, an unconditional standard deviation of about 0.08, while the anchor still absorbs the predictable part of both. The cap at 0.18, near 24% unconditionally, admits perceived risk as large as the realized proportional standard deviation of the after-tax

return to aggregate business capital, 17.7% in post-war US data (Gomme et al., 2011).

- **Wage-Capital return risk correlation** $\rho_{wr,j}$

$$\rho_{wr,j} \sim \text{Unif}(-0.75, 0.75)$$

The correlation between the wage and private-return innovations in (63). The prior is symmetric, reflecting genuine uncertainty about whether subjective wage and return surprises co-move positively or offset one another, and is bounded strictly inside $(-1, 1)$ so that the innovation covariance matrix stays well-conditioned and positive definite for every accepted draw.

- **Forecast horizon** Δ

$$\Delta \sim \text{TPois}(5; 1, 10)$$

The forecast horizon, in years, over which CE anchors are computed before being frozen at their terminal values. The prior places most mass on a roughly five-year horizon, capturing forward-looking but bounded rational agents who cannot credibly forecast a non-stationary transition economy over longer windows.

B.7 Initial Aggregate State

These priors set the aggregate state at each country’s catch-up start $t_{s,j}$ (1982 for China, 1962 for Korea).

- **CN initial state-capital share** $K_{s,cn,1982}/(K_{p,cn,1982} + K_{s,cn,1982})$

$$K_{s,cn,1982}/(K_{p,cn,1982} + K_{s,cn,1982}) \sim \mathcal{TN}(0.6, 0.08^2; 0.5, 0.8)$$

This share is centered above one half because in the 1982 base year state-owned enterprises and public infrastructure dominated the capital stock. Because this is an initial condition that the data identify only weakly—they pin down the total $K_p + K_s$ but not its split—I keep the spread moderate rather than tight. The lower truncation holds the state share above one half, consistent with a still-planned economy, while the upper bound leaves room for the collective and household capital already present.

- **KR initial state-capital share** $K_{s,kr,1962}/(K_{p,kr,1962} + K_{s,kr,1962})$

$$K_{s,kr,1962}/(K_{p,kr,1962} + K_{s,kr,1962}) \sim \mathcal{TN}(0.4, 0.1^2; 0.3, 0.7)$$

Korea’s initial state share is centered well below China’s because private sector already held a sizable part of the capital stock by the early 1960s, before the state-led industrial drive matured. The wider spread acknowledges greater uncertainty about the public-private split in early-1960s Korea, and the bounds keep the share strictly interior so that both capital types enter the gross-complement CES nest (7) with positive mass.

- **Initial endowment-capital ratio** $\mathcal{X}_{\text{start}}/K_{p,\text{start}}$

$$\mathcal{X}_{\text{start}}/K_{p,\text{start}} \sim \mathcal{TN}(1.17, 0.1^2; 1, 1.5)$$

The initial period endows agents with return-scaled wealth (without labor income), aggregate private capital scaled by the data-anchored initial after-tax gross return,

$$\mathcal{X}_{\text{start}} = (1 + (1 - \tau_k)r_p) K_{p,\text{start}},$$

so the ratio to private capital is the after-tax gross return itself, roughly 1.17 at the initial private-capital return of about 20%, and I center the prior there. The lower bound of one is the no-return floor, while the upper bound permits higher capital returns.

Appendix C CE Agent Network Details

This appendix details the sampling of the training problems, the input features, network architectures, loss functions, and optimization of the three CE agent nets of Section 3.3. All three nets consume the same feature vector and share the residual trunk of Figure 29; they differ only in depth and in the task-specific head.

Sampling One training problem is one draw of a CE environment and of an agent inside it. A country, China or Korea with equal probability, and a calendar year, uniform over that country’s macro-data years, come first. The year fixes the policy regime: the consumption-floor break and Korea’s post-1997 capital tax.

The structural parameters and the population-level preference hyperparameters are drawn from the priors of Table 1, and the agent’s own preferences from the population distributions those hyperparameters parameterize. The labor-tax schedule is drawn uniformly over the country’s observed bracket schedules and then perturbed as described below.

The sampled anchor price path, the wage and private-capital return the agent forecasts over, is built from the observed data for that country and year. The return path applies the

sampled income shares, depreciation, and state-capital split to the country’s observed output and capital series, private capital being the drawn split of observed capital. The wage path divides the labor share of output by the effective labor implied by the observed age masses, the cohort education shares, and the sampled human-capital profiles.

Two devices widen the price coverage beyond the observed series: the income shares and average hours follow bounded random walks across the path, and two log-normal disturbances, one for the path level and one per year, perturb it. The path projects flat beyond the last data year, and the agent-side path applies the forecast-horizon freeze and one realized price-risk tilt (61).

The agent is then placed in that environment. Her entry cohort is drawn uniformly from the cohorts alive and not terminal in that year, and her education follows the cohort’s observed shares. Human capital is seeded at the education premium and simulated from schooling exit to the current age through (31) under a random past work intensity, so every sampled age and human-capital pair is jointly reachable; productivity is drawn from the stationary distribution of its AR(1) law.

Initial wealth is income-scaled. A wealth-to-income multiple, drawn from a mixture of a five-percent point mass at exactly zero (covering the zero-bequest entrants the simulator produces), a constrained uniform, a log-normal bulk, and a Pareto tail calibrated to a median near 1.7 and a mean near 3.1 times income, multiplies the agent’s reference income at the frozen anchor wage.

Input Features Each sampled CE problem is encoded into a feature vector of dimension $d = 180$: 34 scalars, a 3-dimensional education indicator, two 70-slot sequences, and 3 state features.

- **Scalars (34).** The preference parameters $(\beta_i, \gamma_i, \psi_i, \nu_i, \log \chi_i)$; the human-capital parameters of (31), $(\delta_h, \gamma_{h,0}, \gamma_{h,1}, \alpha_h, \mathfrak{d}_h)$, and the education premia $(\Delta h_h, \Delta h_c)$; the labor-productivity shock parameters $(\eta_\varepsilon, \sigma_\varepsilon)$; the capital tax τ_k ; an eight-element labor-tax block summarizing the exact bracket schedule, detailed below; the bequest tax τ_m and the consumption-floor ratio Λ ; the price-risk parameters $(\eta_w, \eta_r, \sigma_w, \sigma_r)$ and the realized tilt (ξ_w, ξ_r) ; and the age ratios $(a_0/70, d/70, a_e/70)$.
- **Education (3).** A one-hot indicator of the agent’s education level (below high school, high school, college).
- **Sequences (2×70) .** The log wage path and the private-capital return path, both before

tax, laid out by offset from the current age: slot k holds the value at age $a_0 + k - 1$ over the remaining life and is zero beyond the exit age d .

- **State (3).** $\log(1 + k)$, h , and ε : the predetermined saving, human capital, and labor-productivity shock of (50).

Monetary features are in absolute thousands of 2021 USD and only log-compressed. All features are z-scored with statistics computed on the training split.

Labor-Tax Features The decision problem retains the exact progressive schedule of Appendix A as the labor income tax liability $\mathcal{T}_{l,j,t}$, the tax owed on gross labor income, and the eight-element block above summarizes it by four schedule features (the HSV shape contributing two numbers) and three effective rates. The schedule features are (i) the zero-tax threshold, the lower edge of the first positively-taxed bracket; (ii) the Heathcote–Storesletten–Violante (HSV) (Heathcote et al., 2017) shape (λ_l, τ_l) , net income $\lambda_l y^{1-\tau}$, fit to the schedule by matching the bracket tax in levels and passed as $(\log \lambda_l, \tau_l)$; (iii) the top-rate threshold, the lower edge of the top bracket; and (iv) the top marginal rate. The effective rates are $\mathcal{T}_{l,j,t}(y)/y$ at the three incomes $\{0.3, 0.5, 1\} y_{\text{ref}}$, where y_{ref} is the agent’s average anchor wage over her remaining life, the gross labor income at full efficient labor $q = 1$, so the network sees the realized burden over the income range where agents actually sit (the efficient-labor multiplier $q = h \ell e^\varepsilon$ is typically in $[0.3, 1]$). The exact brackets remain the object that labels net income, so this summary introduces no approximation into the training targets.

To let the surrogates price counterfactual tax policies beyond the historical record, each training draw perturbs the schedule it samples: every bracket edge is scaled by $e^{\sigma_{\text{thr}} z_{\text{thr}}}$ and every marginal rate by $e^{\sigma_{\text{rate}} z_{\text{rate}}}$, with $z_{\text{thr}}, z_{\text{rate}} \sim \mathcal{N}(0, 1)$, $\sigma_{\text{thr}} = 0.25$ and $\sigma_{\text{rate}} = 0.20$ (rates capped at 0.95). Common scaling preserves the monotone bracket structure; the perturbed schedule both labels net income and is re-summarized into all eight features.

Network Architecture Throughout, $\text{Dense}(a \rightarrow b, \sigma)$ denotes a fully connected layer that takes a units in and returns b units out, followed by the element-wise activation σ ; the layer is linear when σ is omitted. An arrow between two layers means the first feeds the second.

The shared trunk of the three nets is a $\text{Dense}(180 \rightarrow 256, \text{swish})$ followed by n_b residual blocks; each block applies a linear $\text{Dense}(256 \rightarrow 256)$ followed by $\text{Dense}(256 \rightarrow 256, \text{swish})$, optionally a dropout layer, and adds the block input back through an identity skip connection, so a block maps $x \mapsto x + \text{swish}(W_2(W_1 x + b_1) + b_2)$. The identity skip keeps gradients well-

conditioned, so the trunk can be made deeper without the vanishing gradients of a plain stack. The heads differ by task:

- **Value net** ($n_b = 3$, no dropout): a linear scalar head $\text{Dense}(256 \rightarrow 1)$ returning the standardized asinh-compressed CE value.
- **Saving net** ($n_b = 4$, dropout = 0.1 in each block): a compression layer $\text{Dense}(256 \rightarrow 128, \text{swish})$ and a path head $\text{Dense}(128 \rightarrow 11)$, one saving rate per offset of the clearing window $[t : t + \Delta]$ with $\Delta \leq 10$. Each rate is clamped to $[0, 1]$, so the net can emit an exact 0 at the borrowing-constrained corner and an exact 1 at the fully-invested corner without a separate saver classifier. Dropout is on because the masked, 11-wide path head is more prone to overfitting than a scalar head.
- **Labor net** ($n_b = 3$, no dropout): a head $\text{Dense}(256 \rightarrow 3, \text{sigmoid})$ returning the normalized working-time parameters $(u_1, u_2, u_3) \in [0, 1]^3$. The labor parameters decode into the inverted-U working-time path. With peak age $p = a_e + u_1(d - a_e)$, peak working time $\ell_p = \text{clamp}(u_2, 0.02, 0.98)$, and curvature $c = 3 + 4u_3$, the path over the working window $[a_e, d]$ is

$$\hat{\ell}_a = \text{sigmoid} \left(\text{logit}(\ell_p) - c \left(\frac{a - p}{d - a_e} \right)^2 \right), \quad a \in [a_e, d], \quad (\text{C.1})$$

the logit-quadratic parameterization used by the Nelder–Mead label search. The three decode ranges are exactly the domain that search optimized over, so the net never proposes a plan outside the domain the solver covered. Bounding the curvature away from zero, $c \in [3, 7]$, keeps the quadratic a realistic life-cycle hump, and flat profiles instead come from ℓ_p near its bounds.

Labor Loss Throughout the three losses, ι indexes the sampled CE problems, and a star marks the exact optimum computed by the PTD solver with the Nelder–Mead search, the label the net regresses on. The labor net is trained in path space, not control space: the three predicted parameters are decoded through (C.1) and compared with the optimal path on the working window,

$$\mathcal{L}_{\text{labor}} = \frac{\sum_{\iota} \sum_{a=a_e^{\iota}}^{d^{\iota}} (\hat{\ell}_a^{\iota} - \ell_a^{*,\iota})^2}{\sum_{\iota} (d^{\iota} - a_e^{\iota} + 1)}. \quad (\text{C.2})$$

The three parameters are weakly identified (different parameters can imply nearly identical paths), so regressing them directly wastes capacity; an A/B comparison showed the path loss cuts the tail regret by about two thirds relative to a parameter-space MSE.

Saving Loss The saving net minimizes the masked MSE of the clamped rate path over each problem’s clearing window,

$$\mathcal{L}_{\text{saving}} = \frac{\sum_{\iota} \sum_{k=0}^{\Delta_{\iota}} (\hat{s}_k^{\iota} - s_k^{*,\iota})^2}{\sum_{\iota} (\Delta_{\iota} + 1)}, \quad (\text{C.3})$$

where offsets beyond a problem’s horizon are masked out of both the numerator and the denominator.

Value Loss The optimal values span several orders of magnitude and both signs, so the value net regresses the standardized asinh-compressed target

$$v^{*,\iota} = \frac{\text{asinh}(V^{*,\iota}/\kappa) - \mu_v}{\sigma_v}, \quad (\text{C.4})$$

by plain MSE, where κ is the median of $|V^*|$ on the training split and (μ_v, σ_v) are the training mean and standard deviation of the compressed values. Predictions are mapped back to value units by inverting (C.4).

Optimization Each net is trained with the AdamW optimizer at a learning rate of 3×10^{-4} , on shuffled mini-batches of 512 problems in single precision. Training runs to a ceiling of 150 epochs with early stopping: it halts once the validation loss has gone 10 consecutive epochs without improving, and the parameters kept are the best-validation checkpoint rather than the final epoch.

Appendix D Rollout Network Details

This appendix details the input features, network architecture, label solver, training procedure, and test metric of the rollout policy net of [Section 3.4](#). Its training states are drawn by the same sampler as the CE agent nets, described in [Appendix C](#), on a disjoint seed range.

Input Features The rollout policy consumes the same feature vector as the CE agent nets, described in [Appendix C](#), encoding the agent’s state and CE environment. It is standardized with the same statistics as the frozen value net that supplies its continuation.

Network Architecture The rollout policy is deliberately lighter than the residual surrogates of [Figure 29](#), because it feeds the per-year, per-agent simulation loop, where forward-pass cost dominates. It is a plain MLP $\text{Dense}(180 \rightarrow 128, \text{swish}) \rightarrow \text{Dense}(128 \rightarrow 128, \text{swish}) \rightarrow \text{Dense}(128 \rightarrow 2, \text{sigmoid})$,

whose two outputs are the current-age controls: the saving rate out of cash-on-hand, clamped to $[0, 0.95]$, and the working time, clamped to $[0.001, 1]$. The save-nothing corner and full-time work are economically valid and numerically regular, so both are admitted. The saving cap keeps consumption away from the CRRA utility divergence at zero consumption, which the consumption floor (48) cannot rule out because it floors cash-on-hand rather than consumption itself, while the labor floor keeps the working time strictly positive.

Label Solver The optimal decisions of the CE problem are not optima of the stochastic rollout objective, so the labels solve the rollout problem itself: for each sampled state, the two controls maximize (66) with the pre-trained CE value net frozen as the continuation and evaluated on the tilted anchor path (65). The double integral is a Gauss–Hermite tensor with three nodes per shock dimension (27 nodes) over the correlated one-step price-risk innovation (ξ_w, ξ_r) of (63) and the labor-productivity innovation of (28).

Each state’s maximization runs in two stages. A batched 13×9 grid over the control bands supplies global coverage; a Nelder–Mead simplex then refines the two best distinct starts, the grid argmax and the near-optimal grid point closest to the CE saving net’s rate, with candidates projected onto the bands inside the objective.²⁵ When the two refined optima agree within 0.1% of the optimal value, the state sits on a plateau of the objective, and the label keeps the candidate whose saving rate is closest to the CE saving net’s decision. Plateaus are the rule, not the exception: 63.5% of states tie-break, which is why an objective-gradient trainer, receiving no signal on exactly those states, previously left the policy unpinned there. States are labeled in chunks of 5,000, each saved as a resumable part file, so an interrupted run loses at most one chunk; labeling 100,000 states takes about 75 minutes on 24 CPU threads.

Training Procedure The label states are drawn by the CE problem sampler of Appendix C on a seed range disjoint from the CE training set and from the held-out test states below. Training is plain supervised regression: mean squared error of the two sigmoid heads against the argmax labels, Adam (learning rate 10^{-3}) on mini-batches of 2,048, a 90/10 train–validation split, and early stopping once ten epochs pass without validation improvement. The run early-stops after 99 epochs at the epoch-89 checkpoint (Figure 32), a few minutes of training. On the validation labels, the saving head’s MAE is 0.10 and the working-time head’s is 0.05; the large saving MAE relative to the optimality gaps below reflects the plateaus, on which label distance

²⁵A gradient-based alternative (box-constrained L-BFGS-B with the objective value and its two-control gradient fused into one forward-mode automatic-differentiation sweep per iterate) reaches the same optima on 93% of benchmark states at no lower cost: with two controls, the simplex needs only cheap objective evaluations and has no slow convergence tail, so the derivative-free refiner is retained.

is cheap in value terms. The labels’ own scatter sets the scale: their mean distance from the CE saving net’s rate is 0.12, so the net sits closer to its targets than the plateau selection is to its tie-break anchor.

Test Metric Because the plateaus make the argmax ill-conditioned, the headline metric is the optimality gap of the objective itself. For each of 5,000 fresh test states, the two-control maximization is re-solved on the identical 27-node objective by a 16×16 grid over the clamp bands followed by three shrinking 9×9 local grids around the running argmax (final resolution about 10^{-3} per control), warm-started at the policy’s own decision, so the oracle value never falls below the net’s. The reported regret is the relative gap $(J^* - \hat{J})/|J^*|$ between the oracle value J^* and the net’s value $\hat{J}$; the decision errors are the median and mean absolute deviations of the two controls from the oracle argmax: on the test states, 0.092 mean and 0.050 median for the saving rate, and 0.046 mean and 0.022 median for the working time.

Appendix E Anchor Network Details

This appendix details the sampling of the training economies, the input features, encoder architecture, training objective, and optimization of the anchor net of [Section 3.5](#), and the two deployment stages that close the residual gap: the on-policy fine-tuning of the deployed net and the learned correction. Unlike the CE agent nets and the rollout policy of [Appendix D](#), which are fitted to solved labels, the initial anchor net is trained without solved labels, directly on the clearing loss.

Sampling One training economy is one draw of a whole auxiliary economy at a single calendar year, carrying $F = 100$ families with one agent per family and alive cohort. The country-year draw, the parameters, and each agent’s education, human capital, productivity and wealth follow the CE problem sampler of [Appendix C](#), with two differences: the demography is the observed one for that year, and the labor-tax schedule is drawn once per economy rather than once per agent, since a common clearing needs a common schedule. Because wealth is drawn agent by agent, the economy’s total wealth follows from its cross-section rather than being set first: aggregate private capital is the population-weighted sum of the individual holdings, and state capital is the drawn share of it. The technology is then calibrated to those aggregates, the CES weights from the drawn factor-share targets and the TFP level from the warm-start return through the firm’s first-order condition, so that the economy starts from a smaller market-clearing loss at its warm start than an uncalibrated draw would. That calibration is approximate

by construction, since the CE labor supply at the anchor is not known when the economy is drawn and an independent hours proxy stands in for it; it is one source of the distribution shift the on-policy fine-tuning below corrects.

Input Features An economy is encoded into two objects: a flat feature vector and the agent cross-section.

- **Flat features (176).** The scalar and short-sequence summary of the economy: the frontier, technology, and nested-CES production parameters; the fiscal block (the consumption-floor, capital-, bequest-, and labor-tax summaries and government efficiency); the aggregate initial state; and, laid out by offset over the clearing window $[t : t + \Delta]$, the data-anchored warm-start wage and return paths together with their warm-start clearing gaps. Offsets beyond a window’s horizon $\Delta + 1 \leq 11$ carry a mask row and are zeroed.
- **Cross-section grid.** The agent pool as an age-by-family tensor: for each of the $H_{\max} = 70$ age slots and each of the $F = 100$ Monte-Carlo families, a 14-dimensional cell holds the family’s log saving, human capital, and productivity state, its one-hot education, its age ratios, and its preference draws; four further per-age channels hold the cohort’s population weight and aggregates. A 70-slot mask marks the alive ages.

Flat features are z-scored, and grid cells and channels are z-scored per channel, with statistics computed on the training split; monetary quantities enter in logs.

Network Architecture The encoder respects the exchangeability of the family pool. A shared Dense($14 \rightarrow 16$, swish) embeds every cell; the F family embeddings at each age are pooled by their mean and standard deviation, so the encoder is invariant to family order and to F . The pooled embedding, concatenated with the four per-age channels, is passed through two width-3 convolutions along the age axis (Conv($36 \rightarrow 32$) $\rightarrow$ Conv($32 \rightarrow 32$), swish, same-padding), which let neighboring cohorts inform one another; a masked global mean-and-max pool over the alive ages collapses the age axis to a 64-dimensional cross-section embedding.

This embedding is concatenated with the 176 flat features and fed to a residual trunk of the same form as the CE agent nets (Figure 29): a Dense($240 \rightarrow 512$, swish) followed by three residual blocks, and a head Dense($512 \rightarrow 22$) returning the wage and return deviations over the eleven window offsets. The head is zero-initialized, so the freshly set-up net emits exactly the data-anchored warm start, and its output is decoded as a bounded multiplicative deviation for

the wage and a bounded logistic deviation for the return, keeping every candidate anchor in an admissible price range.

Training Objective Training minimizes the sampled-economy mean of the imagined-clearing loss (59), computed by a tensorized, single-precision twin of the equilibrium forward pass that clears a full batch of economies on the GPU in one compiled graph. The per-economy loss weights the return residuals so that a five-percentage-point return gap counts like a one-warm-start-wage wage gap, and the batch loss compresses each economy’s loss J logarithmically above a cap of 25, $\rho(J) = \text{cap} \cdot \log(1 + J/\text{cap})$, which is increasing (so each economy’s zero-residual minimum is unchanged) but keeps a few heavy-tailed economies from dominating a mini-batch. The gradient is the exact reverse-mode derivative of $\rho(J)$ through the whole clearing sweep and the frozen agent nets, computed by `Enzyme.jl` through the compiled `Reactant.jl` graph.

Training Procedure One training economy is a full auxiliary-economy draw: a country and a calendar year over its macro-data years, structural and technology parameters from their priors, a clearing horizon Δ , the observed demography in that year, an agent cross-section drawn by the CE problem sampler of Appendix C, and the data-anchored warm-start price path rebuilt from the observed output and capital series. The training pool holds 100,000 such economies, split evenly between the Chinese and Korean priors, drawn and warm-start recalibrated on the GPU, with a held-out validation pool drawn from a disjoint seed range. Each economy carries $F = 100$ families. The net is trained on mini-batches of 16 economies toward a 30,000-step ceiling with a two-phase learning rate under AdamW and a global gradient-norm clip of 1; the run tracks the held-out clearing loss throughout and goes to the ceiling without early stopping, and the saved parameters are its best-validation checkpoint, step 28,800. The full run takes about one hour and forty minutes on the GPU.

On-Policy Fine-Tuning The unsupervised net is trained on sampled economies whose technology is recalibrated to a smaller market-clearing loss at their warm start, but the simulator visits ledger-evolved cross-sections whose warm start is the previous year’s realized prices. This distribution shift is the net’s dominant deployment error: on simulator-visited economies its median relative wage gap is about 20%, and, as the fine-tuning gain below shows, most of that gap is the shift itself rather than a limit of the architecture. I close it by supervised fine-tuning on the deployment distribution. Running the simulator with the exact anchor solve—that is, two damped Gauss–Newton steps—attached, I harvest the encoder inputs and the resulting refined anchor heads at simulator-visited economies as free by-products, 67,409 examples in

eight resumable parts; dropping the degenerate economies whose unscaled clearing loss exceeds a label-hygiene cap of 50, twice the objective’s compression cap, leaves 44,576 pairs (40,113 train, 4,463 validation). The net is fine-tuned on these pairs by minimizing mean squared error over the active wage and return output coordinates. Starting from the unsupervised checkpoint, this fine-tuning cuts the on-policy median relative wage error to 4.6% and the return error from 3.2 to 1.3 percentage points; this sharpened net is the one deployed.

Learned Correction The fine-tuned net still leaves a median on-policy wage error of about 5%. At this feature resolution, the inputs do not capture the clearing operator’s local response, which must be measured rather than inferred from features. Each simulated year is therefore verified and, if needed, corrected. The sharpened anchor is scored by one clearing evaluation, and economies whose root-mean-square window gaps stay within tolerance (5% for the wage, one percentage point for the return) keep their amortized heads. Failing economies are passed through a learned corrector that approximates two damped Gauss–Newton steps and are then re-verified, at a cost of about two clearing evaluations. The correction brings them close to the tolerance rather than always inside it: at the median it lands just past the return margin, as quantified below.

The correction is itself amortized. The exact benchmark consists of two damped Gauss–Newton updates on (59) with a finite-difference Jacobian over the window prices and step damping $\alpha \in \{1, \frac{1}{2}, \frac{1}{4}\}$; one step reaches about 4.4% wage and 0.5 percentage-point return residuals, a second about 0.5% and 0.09 percentage points. Rather than run this survey of finite-difference probes every year, I train a small network, the *corrector*, to emulate the two steps: a multilayer perceptron over the failing economy’s amortized heads, its measured clearing residual, and its flat features (a 220-dimensional input), returning the head correction. Its labels are two-step Gauss–Newton solutions harvested from simulator runs under the production verify-and-correct configuration, so the corrector is trained on exactly the states and failure modes it will face. On held-out gate-failing years the corrector cuts the median distance to that exact solution from 5.6% to 4.7% for the wage and from 1.6 to 1.2 percentage points for the return; its main work is in the tail, where the 90th-percentile return gap falls from 14 to 8 percentage points (Figure E.1).




Figure E.1: Training History of the Corrector Net

Notes: Left: per-step masked mean-squared error of the predicted head correction. Center and right: the corrector’s held-out median relative wage gap and absolute return gap; the dashed line is the uncorrected amortized baseline. The corrector drives both gaps below the baseline and settles within a few thousand steps.

Appendix F Posterior Network Details

This appendix details the summary features, the network architecture and logit-Normal head, the training protocol, and the calibration diagnostics of the posterior net Ψ of [Section 3.6](#).

Input Features The input is the flattened, country-concatenated summary vector $\mathbf{D} \in \mathbb{R}^{752}$: eight statistics per country and year, the capital and output growth rates and the top-1%, top-10%, and bottom-50% pre-tax income and wealth shares, over 37 years for China (1983–2019) and 57 for Korea (1963–2019), giving $8 \times (37 + 57) = 752$ entries. The features are z-scored with statistics computed on the training split. The prior predictive of the growth-rate entries is heavy-tailed, since a few parameter draws produce transiently explosive economies, so the standardized features are clipped at ten standard deviations before entering the network; the clip touches about 0.02% of entries, all in the explosive tail that the observed data never visits, and leaves the network’s inputs well conditioned.

Network Architecture Ψ is a residual multilayer perceptron of the same family as the CE agent nets ([Figure 29](#)): a Dense($752 \rightarrow 256$, swish) input projection, three residual blocks of width 256, and a linear head Dense($256 \rightarrow 136$). The 136 outputs are the location μ_p and log-scale $\log \sigma_p$ of an independent Gaussian for each of the 68 parameters on a transformed scale, with the scale floored below to keep the density proper.

Logit-Normal Posterior Each parameter is estimated on a logit-transformed scale rather than in its native units. Writing $[\ell_p, u_p]$ for a parameter’s prior support and

$y_p = \text{logit}((\theta_p - \ell_p)/(u_p - \ell_p))$ for its logit, the approximate posterior is Gaussian in y_p , so on the native scale it is logit-Normal, $\theta_p = \ell_p + (u_p - \ell_p) \sigma(\mu_p + \sigma_p \varepsilon)$ with $\varepsilon \sim \mathcal{N}(0, 1)$ and $\sigma(\cdot)$ the logistic function. Unlike a truncated-Normal head, the logit-Normal places no mass at the support boundary and needs no boundary correction. The training target is the logit $y_{p,i}$ of each drawn parameter, and the loss is the Gaussian negative log density of $y_{p,i}$ under (μ_p, σ_p) , summed over parameters and averaged over the batch, which is the Monte-Carlo objective (70).²⁶

Training Procedure Because the objective (70) is a Monte-Carlo estimate of the population divergence (68) and not a fit to a fixed target, there is no held-out quantity to select on: more simulated pairs lower the variance of the same integral, and the minimizer is the amortized posterior. I therefore train on all 44,912 non-test pairs, in shuffled mini-batches of 512, with the AdamW optimizer (weight decay 10^{-4}) under a cosine learning-rate schedule annealed from 3×10^{-4} to 10^{-6} and a global gradient-norm clip of 1. In place of validation-based early stopping, the run length is a fixed budget set once by a convergence audit: over a long run the held-out negative log density minimizes early and then degrades as the network begins to memorize the finite simulation sample, so the budget, 100 epochs, is chosen at the audit’s held-out minimum. The full run takes about two minutes on the GPU. Figure F.1 shows the training curve. The loss falls from about 40 at initialization to 18.3 and is still declining slowly at the budget, as the training objective will continue to do; the stopping point comes from the held-out audit above, not from this curve.

²⁶The forecast horizon Δ is integer-valued on $\{1, \dots, 10\}$, so its native endpoints sit exactly on the support bounds and their logits diverge. I pad its transform support by half a step to $[0.5, 10.5]$, so the endpoints map to finite, well-scaled logits; the inverse map is rounded back to the integer horizon, leaving the parameter unchanged.




Figure F.1: Training History of the Posterior Net

Notes: Per-epoch Monte-Carlo negative log density (70) over the 44,912 training pairs. Only the training curve is drawn: the objective is itself an estimate of the expected divergence over the joint prior predictive, so it is the quantity being minimized and has no held-out counterpart to select on.

Calibration Diagnostics I evaluate Ψ on the 4,991-pair held-out test split, where the generating parameters are known because they were drawn. Two frequentist checks of the Bayesian map are reported. *Coverage* is the share of test draws whose true parameter falls inside the posterior’s 90% central interval (Talts et al., 2018). *Held-out contraction*, the ratio of posterior to prior standard deviation averaged over these test datasets, measures how much the summaries identify each parameter, with a value near one flagging a parameter the summaries leave at its prior. It is the prior-predictive counterpart of the contraction reported in Table 1, which divides the posterior standard deviation at the observed data by the same prior standard deviation.

Figure F.2 plots per-parameter coverage and contraction, and Table F.1 summarizes them. Coverage is nominal on average (median 88.3%, with 37 of 68 parameters inside [85%, 95%]), and 28 parameters contract to below 0.9 of their prior spread. The best identified are the education and human-capital technology parameters, the students-to-teachers ratio θ_κ , the college weight increment $\Delta\omega_c$, the learning concavity α_h , and the high-school premium Δh_h , together with the preference means μ_β and μ_γ ; the parameters left near their prior are the household-heterogeneity scale hyperparameters, which the country-level summaries integrate over and so cannot pin down. Because these are averages over the held-out split, a parameter can rank differently here than under the observed-data contraction of Table 1: the forecast horizon Δ contracts to 0.73 on this split but sits at 0.98 at the observed data. The miscalibration sits in the opposite tail: the most sharply contracted parameters (θ_κ , μ_β , $\Delta\omega_c$, α_h) have coverage between 46% and 58%,

well located but overconfidently tight, and China’s subjective price-risk correlation ρ_{wr} is the severe case, driven to a tight but frequently misplaced posterior (coverage about 16%).




Figure F.2: Posterior-Net Calibration and Identification

Notes: Top: the held-out 90%-interval coverage per parameter; the dashed line is the nominal 90%. Bottom: the contraction, the ratio of posterior to prior standard deviation, per parameter; the dashed line is one (no identification). Diagnostics are computed on the 4,991-pair held-out test split.

Table F.1: Posterior-Net Held-Out Calibration

Diagnostic (median over 68 parameters)	Value
90%-interval coverage	88.3%
Parameters with coverage in [85%, 95%]	37/68
Held-out contraction (posterior sd / prior sd)	0.95
Parameters with held-out contraction < 0.9	28/68

Notes: Held-out diagnostics on the 4,991-pair test split. Coverage is the empirical frequency that the true parameter falls in the posterior’s 90% central interval; held-out contraction is the ratio of posterior to prior standard deviation, averaged over the test datasets rather than taken at the observed data as in Table 1.

Appendix G Computation Time

This appendix reports the wall-clock cost of the estimation pipeline. All computations run on a single laptop with 24 CPU threads and one NVIDIA GeForce RTX 5090 Laptop GPU (24 GB), the machine of Section 3. The code is written in Julia: the neural networks are built in Lux.jl, compiled to the GPU with Reactant.jl (XLA), and differentiated with Enzyme.jl, while the CE agents’ and the rollout policy’s training labels are generated on CPU threads. Table G.1 collects the timing of each learned stage, separating the cost of generating its training data from the cost of the training itself.

Table G.1: Computation Time by Stage

Stage	Data Size	Data Gen.	Training
CE agent networks	1,000,000	2 hour 5 min	18 min
Rollout network	100,000	1 hour 15 min	5 min
Anchor network	100,000	–	1 hour 40 min
Anchor fine-tuning	44,576	2 hour 50 min	35 min
Anchor corrector	153,805	2 hour 15 min	2 min
Posterior network (NPE)	49,903	10 hour 50 min	2 min
Total		19 hour 15 min	2 hour 42 min

Notes: Data Gen. is the time to generate a stage’s training data; Training is the network training or fine-tuning time. Data Size is the labeled data a stage generates, before the split into training and held-out parts.

References

- ADAM, K., AND A. MARCET (2011): “Internal rationality, imperfect market knowledge and asset prices,” *Journal of Economic Theory*, 146 (3), 1224–1252.
- AIYAGARI, S. R. (1994): “Uninsured idiosyncratic risk and aggregate saving,” *The Quarterly Journal of Economics*, 109 (3), 659–684.
- AMOS, B. (2023): “Tutorial on amortized optimization,” *Foundations and Trends in Machine Learning*, 16 (5), 592–732.
- AMSDEN, A. H. (1989): *Asia’s Next Giant: South Korea and Late Industrialization*, New York: Oxford University Press.
- AUCLERT, A., B. BARDÓCZY, M. ROGNLIE, AND L. STRAUB (2021): “Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models,” *Econometrica*, 89 (5), 2375–2408.
- AZINOVIC, M., L. GAEGAUF, AND S. SCHEIDEGGER (2022): “Deep equilibrium nets,” *International Economic Review*, 63 (4), 1471–1525.
- BACHAS, P., M. FISHER-POST, A. JENSEN, AND G. ZUCMAN (2026): “Capital Taxation, Development, and Globalization: Evidence from a Macro-Historical Database,” *American Economic Journal: Applied Economics*, Forthcoming.
- BARRO, R. J. (2013): “Education and economic growth,” *Annals of Economics and Finance*, 14 (2), 301–328.
- BARRO, R. J., AND J. W. LEE (2013): “A new data set of educational attainment in the world, 1950–2010,” *Journal of Development Economics*, 104, 184–198.

- BENHABIB, J., AND M. M. SPIEGEL (1994): “The Role of Human Capital in Economic Development: Evidence from Aggregate Cross-Country Data,” *Journal of Monetary Economics*, 34 (2), 143–173.
- BERTSEKAS, D. (2025): *A course in reinforcement learning*: Athena Scientific, 2nd edition.
- BERTSEKAS, D. P., AND J. N. TSITSIKLIS (1996): *Neuro-Dynamic Programming*, Belmont, MA: Athena Scientific.
- BEWLEY, T. F. (1981): “Stationary equilibrium,” *Journal of Economic Theory*, 24 (2), 265–295.
- BLOOM, D., D. CANNING, AND J. SEVILLA (2003): *The demographic dividend: A new perspective on the economic consequences of population change*: Rand Corporation.
- BLOOM, D. E., AND J. G. WILLIAMSON (1998): “Demographic transitions and economic miracles in emerging Asia,” *The World Bank Economic Review*, 12 (3), 419–455.
- BOHK-EWALD, C., P. LI, AND M. MYRSKYLÄ (2018): “Forecast accuracy hardly improves with method complexity when completing cohort fertility,” *Proceedings of the National Academy of Sciences*, 115 (37), 9187–9192.
- CAI, Y., AND K. L. JUDD (2023): “A simple but powerful simulated certainty equivalent approximation method for dynamic stochastic problems,” *Quantitative Economics*, 14 (2), 651–687.
- CAO, D., K. GILLINGHAM, F. ISKHAKOV, A. MUNK-NIELSEN, J. RUST, AND B. SCHJERNING (2026): “Making Waves: The Effect of Macro Shocks on the Auto Market,” Working Paper.
- CHANCEL, L., T. PIKETTY, E. SAEZ, AND G. ZUCMAN (2022): “World Inequality Report 2022,” Technical report, World Inequality Lab, <https://wir2022.wid.world/>.
- CHEN, Z., Z. HE, AND C. LIU (2020): “The financing of local government in China: Stimulus loan wanes and shadow banking waxes,” *Journal of Financial Economics*, 137 (1), 42–71.
- CHIRINKO, R. S. (2008): “ σ : The long and short of it,” *Journal of Macroeconomics*, 30 (2), 671–686.
- COOLEY, T. F., AND E. C. PRESCOTT (1995): “Economic growth and business cycles,” in *Frontiers of Business Cycle Research* ed. by Cooley, T. F. Princeton, NJ: Princeton University Press, Chap. 1, 1–38.

- CRANMER, K., J. BREHMER, AND G. LOUPPE (2020): “The frontier of simulation-based inference,” *Proceedings of the National Academy of Sciences*, 117 (48), 30055–30062.
- DE NARDI, M., E. FRENCH, AND J. B. JONES (2010): “Why do the elderly save? The role of medical expenses,” *Journal of Political Economy*, 118 (1), 39–75.
- FAN, X., A. SESHADRI, AND C. TABER (2024): “Estimation of a life-cycle model with human capital, labor supply, and retirement,” *Journal of Political Economy*, 132 (1), 48–95.
- FANG, H., AND R. ZHANG (2025): “Corruption stereotype and the unintended consequences of an anti-corruption campaign: evidence from the real estate sector in China,” *Journal of Public Economics*, 249, 105474.
- FEENSTRA, R. C., R. INKLAAR, AND M. P. TIMMER (2015): “The next generation of the Penn World Table,” *American Economic Review*, 105 (10), 3150–3182.
- FERNÁNDEZ-VILLAYERDE, J. (2026): “Deep Learning for Solving Economic Models,” *Journal of Economic Literature*, Forthcoming.
- FERNÁNDEZ-VILLAYERDE, J., S. HURTADO, AND G. NUÑO (2023a): “Financial frictions and the wealth distribution,” *Econometrica*, 91 (3), 869–901.
- FERNÁNDEZ-VILLAYERDE, J., L. E. OHANIAN, AND W. YAO (2023b): “The neoclassical growth of China,” working paper, NBER.
- FLODÉN, M., AND J. LINDÉ (2001): “Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance?” *Review of Economic Dynamics*, 4 (2), 406–437.
- FRENCH, E., AND J. B. JONES (2011): “The effects of health insurance and self-insurance on retirement behavior,” *Econometrica*, 79 (3), 693–732.
- GECHERT, S., T. HAVRANEK, Z. IRSOVA, AND D. KOLCUNOVA (2022): “Measuring capital-labor substitution: The importance of method choices and publication biases,” *Review of Economic Dynamics*, 45, 55–82.
- GILLINGHAM, K., F. ISKHAKOV, A. MUNK-NIELSEN, J. RUST, AND B. SCHJERNING (2022): “Equilibrium trade in automobiles,” *Journal of Political Economy*, 130 (10), 2534–2593.
- GOMME, P., B. RAVIKUMAR, AND P. RUPERT (2011): “The return to capital and the business cycle,” *Review of Economic Dynamics*, 14 (2), 262–278.

- GOURIÉROUX, C., AND A. MONFORT (1996): *Simulation-Based Econometric Methods*: Oxford University Press.
- GRANDMONT, J. M. (1977): “Temporary general equilibrium theory,” *Econometrica*, 535–572.
- HA, D., AND J. SCHMIDHUBER (2018): “Recurrent World Models Facilitate Policy Evolution,” in *Advances in Neural Information Processing Systems* Volume 31.
- HAGGARD, S. (2018): *Developmental states*: Cambridge University Press.
- HAN, J., Y. YANG, AND W. E (2026): “DeepHAM: A global solution method for heterogeneous agent models with aggregate shocks,” *Quantitative Economics*, 17 (2), 297–341.
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2017): “Optimal Tax Progressivity: An Analytical Framework,” *The Quarterly Journal of Economics*, 132 (4), 1693–1754.
- HERBST, E. P., AND F. SCHORFHEIDE (2015): *Bayesian estimation of DSGE models*: Princeton University Press.
- HONG, S., N. N. KIM, Z. MO, AND L. YANG (2025): “Income inequality in South Korea, 1933–2022: Evidence from distributional national accounts,” working paper, World Inequality Lab.
- HUBBARD, R. G., J. SKINNER, AND S. P. ZELDES (1995): “Precautionary saving and social insurance,” *Journal of Political Economy*, 103 (2), 360–399.
- HUSSAIN, A. (1994): “Social security in present-day China and its reform,” *American Economic Review*, 84 (2), 276–280.
- HWANG, H., A. PURWIN, AND J. PARELIUSSEN (2022): “Strengthening the Social Safety Net in Korea,” OECD Economics Department Working Papers 1733, OECD Publishing, Paris.
- IMAI, S., AND M. P. KEANE (2004): “Intertemporal labor supply and human capital accumulation,” *International Economic Review*, 45 (2), 601–641.
- ISLAH, B., AND B. LIGHT (2025): “Bounded Foresight Equilibrium in Large Dynamic Economies with Heterogeneous Agents and Aggregate Shocks,” *arXiv preprint arXiv:2502.16536*.
- JOHNSON, C. (1982): *MITI and the Japanese miracle: the growth of industrial policy, 1925–1975*: Stanford university press.

- JUDD, K. L., L. MALIAR, AND S. MALIAR (2011): “Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models,” *Quantitative Economics*, 2 (2), 173–210.
- KAPLAN, G., K. MITMAN, AND G. L. VIOLANTE (2020): “The housing boom and bust: Model meets evidence,” *Journal of Political Economy*, 128 (9), 3285–3345.
- KARABARBOUNIS, L., AND B. NEIMAN (2014): “The global decline of the labor share,” *The Quarterly Journal of Economics*, 129 (1), 61–103.
- KASE, H., L. MELOSI, AND M. ROTTNER (2025): “Estimating Nonlinear Heterogeneous Agent Models with Neural Networks,” BIS Working Paper 1241, Bank for International Settlements.
- KLENKE, A. (2020): *Probability theory: a comprehensive course*: Springer, 3rd edition.
- KOH, Y. (2015): “The social safety net in Korea: from welfare to workfare,” in *Social Policies in an Age of Austerity*: Edward Elgar Publishing, 44–80.
- KOHLI, A. (2004): *State-directed development: Political power and industrialization in the global periphery*: Cambridge University Press.
- KRUEGER, D., K. MITMAN, AND F. PERRI (2016): “Macroeconomics and household heterogeneity,” in *Handbook of Macroeconomics* Volume 2: Elsevier, 843–921.
- KRUGMAN, P. (1994): “The Myth of Asia’s Miracle,” *Foreign Affairs*, 73 (6), 62–78.
- KRUSELL, P., AND A. A. SMITH, JR (1998): “Income and wealth heterogeneity in the macroeconomy,” *Journal of Political Economy*, 106 (5), 867–896.
- KUBLER, F., AND S. SCHEIDEGGER (2025): “Self-Justified Equilibria: Existence and Computation,” *Journal of the European Economic Association*.
- LANE, N. (2025): “Manufacturing revolutions: Industrial policy and industrialization in South Korea,” *The Quarterly Journal of Economics*, 140 (3), 1683–1741.
- LI, H., L. LI, B. WU, AND Y. XIONG (2012): “The end of cheap Chinese labor,” *Journal of Economic Perspectives*, 26 (4), 57–74.
- LI, R., H. SONG, J. ZHANG, AND J. ZHANG (2025): “The effects of a multifaceted poverty alleviation program on rural income and household behavior in China,” *American Economic Journal: Economic Policy*, 17 (2), 319–357.

- LIM, W. (2003): “Public Enterprise Reform and Privatization in Korea: Lessons for Developing Countries,” KDI Policy Study, Korea Development Institute (KDI).
- LIN, J. Y. (1992): “Rural reforms and agricultural growth in China,” *American Economic Review*, 34–51.
- LIU, W., AND F. ISKHAKOV (2025): “Policy Tree Descendant Algorithm for Solving Consumption-Savings Problems with Safety Net,” working paper, Australian National University.
- LUCAS, R. E., JR. (1988): “On the Mechanics of Economic Development,” *Journal of Monetary Economics*, 22 (1), 3–42.
- LUCAS, R. E., JR. (1993): “Making a Miracle,” *Econometrica*, 61 (2), 251–272.
- MALIAR, L., AND S. MALIAR (2022): “Deep learning classification: Modeling discrete labor choice,” *Journal of Economic Dynamics and Control*, 135, 104295.
- MALIAR, L., S. MALIAR, AND P. WINANT (2021): “Deep learning for solving dynamic economic models,” *Journal of Monetary Economics*, 122, 76–101.
- MCLEATH, R. (2020): *Statistical rethinking: A Bayesian course with examples in R and Stan*: Chapman and Hall/CRC, 2nd edition.
- MCFADDEN, D. (1989): “A method of simulated moments for estimation of discrete response models without numerical integration,” *Econometrica*, 57 (5), 995–1026.
- MOLL, B. (2026): “The trouble with rational expectations in heterogeneous agent models: A challenge for macroeconomics,” *The Economic Journal*, 136 (676), 1173–1216.
- MURPHY, K. P. (2023): *Probabilistic Machine Learning: Advanced Topics*: MIT Press.
- NELSON, R. R., AND E. S. PHELPS (1966): “Investment in Humans, Technological Diffusion, and Economic Growth,” *American Economic Review*, 56 (1/2), 69–75.
- PAPAMAKARIOS, G., AND I. MURRAY (2016): “Fast ε -free Inference of Simulation Models with Bayesian Conditional Density Estimation,” in *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, 1028–1036.
- PEARL, J. (2009): *Causality: Models, Reasoning, and Inference*: Cambridge University Press, 2nd edition.

- PEARL, J., M. GLYMOUR, AND N. P. JEWELL (2016): *Causal Inference in Statistics: A Primer*: John Wiley & Sons.
- PIKETTY, T., L. YANG, AND G. ZUCMAN (2019): “Capital accumulation, private property, and rising inequality in China, 1978–2015,” *American Economic Review*, 109 (7), 2469–2496.
- RUST, J. (1987): “Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher,” *Econometrica*, 999–1033.
- SILVER, D., A. HUANG, C. J. MADDISON, A. GUEZ, L. SIFRE, G. VAN DEN DRIESSCHE, J. SCHRITTWIESER ET AL. (2016): “Mastering the game of Go with deep neural networks and tree search,” *Nature*, 529 (7587), 484–489.
- SIMON, H. A. (1955): “A behavioral model of rational choice,” *The Quarterly Journal of Economics*, 99–118.
- SIMON, H. A. (1996): *The Sciences of the Artificial*: MIT Press, 3rd edition.
- SONG, Z. M., AND W. XIONG (2024): “The Mandarin Model of Growth,” Working paper, October 2024.
- STACHURSKI, J., AND A. A. TODA (2019): “An impossibility theorem for wealth in heterogeneous-agent models with limited heterogeneity,” *Journal of Economic Theory*, 182, 1–24.
- TALTS, S., M. BETANCOURT, D. SIMPSON, A. VEHTARI, AND A. GELMAN (2018): “Validating Bayesian inference algorithms with simulation-based calibration,” *arXiv preprint arXiv:1804.06788*.
- VANDENBUSSCHE, J., P. AGHION, AND C. MEGHIR (2006): “Growth, Distance to Frontier and Composition of Human Capital,” *Journal of Economic Growth*, 11 (2), 97–127.
- WADE, R. (1990): *Governing the Market: Economic Theory and the Role of Government in East Asian Industrialization*, Princeton, NJ: Princeton University Press.
- WALDER, A. G. (2015): *China under Mao: A revolution derailed*: Harvard University Press.
- WOODFORD, M. (2019): “Monetary Policy Analysis When Planning Horizons Are Finite,” *NBER Macroeconomics Annual*, 33, 1–50.
- YANG, D. T. (1999): “Urban-biased policies and rising income inequality in China,” *The American Economic Review*, 89 (2), 306–310.

- YANG, Y., C. WANG, A. SCHAAB, AND B. MOLL (2026): “Structural Reinforcement Learning for Heterogeneous Agent Macroeconomics,” Working paper, first version December 2025.
- YOO, I. (2000): “Experience with tax reform in the Republic of Korea,” *Asia Pacific Development Journal*, 7 (2), 75–104.
- YOUNG, A. (1995): “The tyranny of numbers: confronting the statistical realities of the East Asian growth experience,” *The Quarterly Journal of Economics*, 110 (3), 641–680.
- YOUNG, A. (2003): “Gold into base metals: Productivity growth in the People’s Republic of China during the reform period,” *Journal of Political Economy*, 111 (6), 1220–1261.
- ZAMMIT-MANGION, A., M. SAINSBURY-DALE, AND R. HUSER (2024): “Neural methods for amortized inference,” *Annual Review of Statistics and Its Application*, 12.
- ZHANG, H., H. ZHANG, AND J. ZHANG (2015): “Demographic age structure and economic development: Evidence from Chinese provinces,” *Journal of Comparative Economics*, 43 (1), 170–185.
- ZHOU, Y., X. HUANG, Y. SHEN, AND L. TIAN (2023): “Does targeted poverty alleviation policy lead to happy life? Evidence from rural China,” *China Economic Review*, 81, 102037.
- ZHU, X. (2012): “Understanding China’s growth: Past, present, and future,” *Journal of Economic Perspectives*, 26 (4), 103–124.