Agent-Based Macroeconomic Modeling: Emergent Inequality and Business Cycles from Simple Behavioral Imitation Rules

Mia Zhou
Tsinghua University
Summer Research Internship

Abstract

This report documents a summer research internship at Tsinghua University, in which an Agent-Based Model (ABM) was developed to study the emergence of macroeconomic phenomena — specifically business cycles and wealth inequality — from simple social imitation behaviors. Rather than assuming fully rational households, the model populates an economy with N boundedly rational agents who follow a single behavioral rule: at Poisson-distributed random times, each household observes its social network neighbors and copies the savings rate of whichever neighbor currently displays the highest consumption, plus a small exploration noise. The model produces two key results without any external shocks: (1) a phase transition at a critical social update interval τc, below which the economy falls into a poverty trap, and above which endogenous business cycles and bimodal wealth inequality emerge spontaneously; and (2) aggregate savings rates that approach the Golden Rule optimum despite individual myopia. The model was subsequently extended with LLM-based decision modules, revealing how homophily generates economic echo chambers that alter cycle length and volatility.

Keywords: agent-based model, bounded rationality, business cycles, wealth inequality, social imitation, RCK model, endogenous dynamics, homophily.

1. PROJECT OVERVIEW

This report documents the design, implementation, and findings of an agent-based macroeconomic model built during a summer research internship at Tsinghua University. The central research question is: can the defining features of modern economies — recurring boom–bust cycles and persistent wealth stratification — arise purely from decentralized local imitation, without rational optimization or exogenous shocks?

The model replaces the standard representative rational household of the Ramsey–Cass–Koopmans (RCK) framework with N heterogeneous, myopic agents connected in a social network.

The project proceeded in two phases: a deterministic rule-based ABM, followed by an LLM-based extension to study belief dynamics and echo chambers.

The core thesis is that business cycles and wealth inequality can emerge endogenously from a single behavioral mechanism — boundedly rational households imitating the highest-consuming neighbor in their social network — without any need for exogenous shocks or full rationality.

2. THEORETICAL BACKGROUND

2.1 The RCK Model and Its Limitations

The Ramsey–Cass–Koopmans (RCK) model is the canonical framework for macroeconomic growth theory. In the standard formulation, a single representative household rationally maximizes long-run discounted utility over an infinite horizon. Households earn income from both labor and capital, and the model yields an optimal intertemporal consumption path via dynamic programming. Business cycles, in this framework, are attributed entirely to exogenous shocks — technology shocks in real business cycle models, or demand shocks in New Keynesian models.

However, empirical evidence consistently shows that real households deviate from optimal consumption paths by 30–50%, and intergenerational learning fails to fully correct this gap. This motivates replacing full rationality with a plausible behavioral rule.

2.2 The Golden Rule Benchmark

The optimization benchmark retained from RCK theory is the Golden Rule savings rate, which maximizes long-run steady-state consumption:

$$s_{gold} = 0.5 \quad \Rightarrow \quad Y^* = \frac{s_{gold}}{\delta} = 10, \quad C^* = 5$$

A key finding of this project is that the behavioral ABM collectively approaches this benchmark at the aggregate level, even though no individual agent performs any optimization.

3. MODEL DESIGN AND TECHNICAL ARCHITECTURE

3.1 Agent Setup

The model simulates N heterogeneous households, each endowed with individual capital K_i, a dynamically updated savings rate s_i(t), and income derived from a Cobb–Douglas production function:

$$Y = K^{\alpha}L^{1-\alpha}$$

where α is the capital output elasticity, K is aggregate capital, and L is total labor. Each household's income combines capital returns and wages:

$$I_i = rK_i + w\frac{L}{N}$$

Capital accumulates and depreciates according to:

$$K_i = s_iI_i - \delta K_i = (rs_i - \delta)K_i + ws_i\frac{L}{N}$$

The steady state gives the long-run capital stock K_i* for each savings rate. The capital return rate satisfies r ∝ k^(α−1), which is the key feedback mechanism driving endogenous cycles.

3.2 The Social Imitation Rule

Instead of solving an optimization problem, each household follows a simple behavioral rule at Poisson-distributed random times with mean interval τ. At each update, the household identifies the neighbor j with the highest current consumption and copies that savings rate plus exploration noise ε (±1%):

$$s_i^{new} = s_{argmax}\, j \in \mathcal{N}(i) \quad (C_j) + \varepsilon$$

3.3 Network Structure

Social connections are modeled using an Erdős–Rényi random graph, where each pair of households is connected with probability p, giving average degree ⟨k⟩ = Np and average shortest path length χ. The network topology directly determines the critical threshold τ_c.

3.4 Simulation Pipeline

A strict tick order was enforced to ensure reproducibility:

- Environment updates (compute aggregate K, Y, r, w)
- Agents observe (each household reads neighbor states)
- Agents decide (apply imitation rule if Poisson update is due)
- Actions execute (update s_i)
- Aggregates update (recompute Y, K, Gini coefficient, etc.)
- Metrics log (record to MLflow)

4. KEY FINDINGS

4.1 The Critical Threshold τ_c

The central control parameter is τ, the average time between updates. It maps directly onto the RCK discount rate ρ through:

$$\rho(\tau) = \frac{\frac{\delta}{2}}{e^{\delta\tau/2} - 1}$$

As $\tau \to 0$, households behave as if $\rho \to \infty$ (extremely short-sighted). As $\tau \to \infty$, they approach fully patient long-run optimization ($\rho \to 0$). The endogenous optimal savings rate is:

$$s^*(\tau) = \frac{1 - e^{-\delta\tau/2}}{2 - e^{-\delta\tau/2}}$$

This increases monotonically with τ and approaches $s_{gold} = 0.5$ as $\tau \to \infty$. A critical threshold $\tau_c \approx 250$ separates two qualitatively distinct economic regimes.

4.1.1 Regime 1: Frequent Updates ($\tau < \tau_c$) — Poverty Trap

When agents imitate too frequently, they are overreactive to short-term consumption differences. Since high consumption often reflects low savings, the system drifts toward underinvestment:

- Savings rates collapse toward zero across the population.
- Capital fails to accumulate; GDP remains chronically low.
- Wealth distribution is unimodal — everyone is equally poor.

4.1.2 Regime 2: Infrequent Updates ($\tau > \tau_c$) — Endogenous Cycles and Inequality

When agents update slowly enough, capital accumulation diverges across households. Delayed feedback creates a self-sustaining cycle:

- Wealth inequality: the savings rate distribution becomes bimodal — a distinct rich class (high s, high K) and poor class (low s, low K) emerge from internal dynamics alone.
- Endogenous business cycles: output, capital, and savings rates oscillate persistently without any external shock.
- Near-optimal aggregate performance: the average savings rate approaches $s_{gold} = 0.5$ despite individual myopia.

4.2 Mechanism of the Endogenous Business Cycle

The cycle operates through the interaction between capital returns and imitative behavior, producing a spontaneous limit cycle in (k, c) phase space:

Recovery phase (t_1): Most households are poor workers with low capital and savings. A few capitalists hold high capital and high savings. When poor households update, they imitate the high-saving capitalists and begin accumulating capital. The economy expands.

Peak phase (t_2): Many households have switched to high savings and accumulated capital. The return on capital r falls. Rich households now observe poor households with low savings but high short-run consumption, and mistakenly imitate them — reducing savings and depleting capital. The economy contracts.

This cycle is purely endogenous — it requires no external shock. A key empirical prediction: savings rates rise before output recovers from recession, consistent with observed macroeconomic data.

4.3 Network Structure and the Critical Threshold

The critical threshold τ_c is determined by network topology according to the power law:

$$\tau_c \sim \frac{e^{-\chi}}{\langle k \rangle}$$

- ⟨k⟩ (average degree): more neighbors to observe means behavioral shifts spread faster, stabilizing the system.
- χ (average shortest path): denser networks (smaller χ) allow faster propagation of imitation, further stabilizing the system.

This relationship was verified across network sizes N = 100 to N = 400. Larger systems are also more stable: more agents imply smaller per-agent deviations $\Delta = \tilde{s} - s^*$, reducing the likelihood of crossing τ_c.

4.4 Matching Real Business Cycle Timescales

Post-WWII G7 economies exhibit business cycles of approximately 9–20 years. Using a Barabási–Albert scale-free network with 4,000 households, average degree 40, δ = 0.2, and τ = 1.5:

- Pure imitation model: cycle ≈ 28 years.
- Adding 5% rational households (fixed at $s^* = 2/9$): cycle ≈ 10.5 years — within the empirical range.

Homophily further modulates cycle length: networks where agents prefer similar neighbors slow behavioral diffusion, extending cycles (e.g., 11 → 17 years), while random connections shorten them.

5. LLM EXTENSION: BELIEF DYNAMICS AND ECHO CHAMBERS

5.1 Motivation and Framework

After establishing the rule-based behavioral baseline, the model was extended by replacing the deterministic imitation rule with LLM-based decision modules, using the AgentSociety framework to manage simulation environment and agent interactions. The goal was to test whether autonomous agents with richer internal reasoning could replicate or complexify the patterns observed under simple imitation.

5.2 Implementation

Key design choices ensured comparability against the rule-based baseline:

- Structured outputs: the LLM was constrained to output JSON-formatted actions given explicit state inputs and risk constraints.
- State inputs: each agent received its own capital, income, neighbor states, and aggregate economic indicators at each decision step.
- Experiment tracking: all runs were logged with MLflow, including random seeds, network topology, and LLM parameters, ensuring full reproducibility.

5.3 Belief Dynamics Under Homophily

The most significant finding from the LLM extension concerned homophily — the tendency of agents to interact preferentially with similar neighbors:

- Without homophily: belief distributions remained unimodal and relatively stable, consistent with the rule-based model.
- With strong homophily: belief distributions shifted from unimodal to bimodal. Distinct echo chambers formed, with clusters reinforcing each other's beliefs while diverging from other clusters.

Polarization was quantified using variance, cluster separation, and distribution shape over time. Echo chambers significantly altered economic cycle volatility and length, consistent with the network structure findings from the rule-based model.

6. ROBUSTNESS AND REPRODUCIBILITY

To ensure the reliability of all findings:

- Random seeds were fixed across all simulation runs.
- Configurations were versioned and stored alongside experiment logs.

- MLflow was used for experiment tracking, logging all parameter choices, network topologies, and metric time series.
- Sensitivity analysis was performed over imitation strength, exploration noise (ε), and network connectivity (p, N).
- The deterministic rule-based baseline was validated independently before introducing any stochastic or LLM components.

7. DISCUSSION

7.1 Challenging Standard Macroeconomic Theory

Standard macroeconomic models require either fully rational agents or exogenous shocks to generate realistic dynamics. This project demonstrates that neither is necessary. Business cycles can arise entirely from decentralized social imitation, wealth inequality can emerge spontaneously from a symmetric initial distribution, and near-optimal aggregate performance can be achieved by myopic agents.

7.2 Policy Implications

The model suggests several policy-relevant insights. First, update frequency matters: policies that moderate the speed of social financial imitation could help economies avoid poverty traps. Second, network density matters: denser, better-connected social networks are more stable and less prone to inequality-generating phase transitions. Third, a small rational anchor is sufficient: a 5% fraction of agents maintaining long-run savings targets significantly shortens and regularizes business cycles.

7.3 LLMs as a Tool for ABM Research

The LLM extension demonstrates a productive research methodology: use simple deterministic rules to establish baseline dynamics, then introduce LLM agents to explore richer behavioral hypotheses. Constraining LLMs to structured outputs is critical for maintaining experimental comparability.

8. CONCLUSION

This internship produced an agent-based macroeconomic model demonstrating that business cycles and wealth inequality can be jointly explained by a single behavioral mechanism: myopic social imitation. The key contributions are:

- A phase transition from stable poverty traps to oscillatory inequality, governed by the single control parameter τ.

- A fully endogenous business cycle mechanism requiring no external shock, sustained by individual-level randomness alone.
- A direct analytical mapping between behavioral myopia and the classical RCK discount rate ρ.
- Cycle lengths consistent with empirical G7 data, particularly with a small fraction of rational agents.
- Network structure — size, density, and homophily — as a quantifiable determinant of macroeconomic stability.
- An LLM-based extension showing how homophily generates economic echo chambers that alter cycle dynamics.

By challenging the twin pillars of traditional macroeconomics — full rationality and exogenous shocks — this work opens a productive research agenda connecting behavioral economics, complex systems, network science, and macroeconomic theory.

References

Ramsey, F. P. (1928). A mathematical theory of saving. Economic Journal, 38(152), 543–559.

Cass, D. (1965). Optimum growth in an aggregative model of capital accumulation. Review of Economic Studies, 32(3), 233–240.

Koopmans, T. C. (1965). On the concept of optimal economic growth. Pontificiae Academiae Scientiarum Scripta Varia, 28, 225–300.

Tesfatsion, L., & Judd, K. L. (Eds.). (2006). Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics. Elsevier.

Erdős, P., & Rényi, A. (1959). On random graphs. Publicationes Mathematicae, 6, 290–297.

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. Science, 286(5439), 509–512.