

The Origins of the Wealth Threshold: Separating Preference Curvature from Financial Frictions

Zihe Wang

Duke Kunshan University, Suzhou, China
zihe.wang@dukekunshan.edu.cn

Abstract. The paper investigates whether wealth accumulation exhibits an endogenous threshold beyond which the marginal wealth–growth relationship structurally moderates. A PPO agent maximizing log utility generates a simulated panel under wealth-dependent borrowing costs. Hansen's panel threshold regression identifies a statistically significant structural break at wealth level 8.17 ($F = 213.5$, $p < 0.001$; 95% CI: [7.60, 8.69]): below the threshold, each log-point increase in wealth reduces the growth rate more steeply; above it, the growth penalty eases. To find out the origin of this breakpoint, the paper conducts two ablation exercises. Firstly, remove borrowing-cost asymmetry yields a threshold of 7.78, while replacing log utility with CRRA preferences ($\gamma=0.5$) gives 7.96. It comes out that all three thresholds cluster tightly around a wealth level of 8. This result illustrates that applying structural break analysis to optimal policy simulations is an effective way to unravel out the behavioral and institutional drivers of wealth inequality.

Keywords: Wealth threshold, Reinforcement learning, Panel Threshold Regression, Preference curvature, Financial frictions

1. Introduction

The intuition that "wealth begets wealth" has long shaped public discourse and policy debates on inequality. In macroeconomic terms, this idea is often framed as a question of acceleration: does an individual's ability to accumulate wealth improve once a certain level of affluence is reached. Piketty's $r > g$ hypothesis formalizes one version of this claim [1], arguing that the rate of return on capital tends to exceed the growth rate of the economy, concentrating wealth at the top. While the $r > g$ logic operates at the aggregate level, its micro-foundation requires understanding how an individual's wealth-growth profile depends on her own wealth level—and whether that profile contains a structural break. Recent empirical work has revisited this hypothesis using long-run historical data: Abatamarco, Lagomarsino and Spiganti propose a non-parametric strategy to test the "third law" of Piketty [1, 2], finding that the $r-g$ gap explains approximately 10% of the trend in U.S. income inequality, while Jakurti [3], using data from 16 advanced economies over 1870–2020, shows that a one-percentage-point increase in the $r-g$ gap raises the top 1% wealth share by 3.7% on average. Yet, it is still a question whether this threshold exists, and if it does, whether it arises from financial asymmetries where the rich borrow more cheaply, the shape of preferences referring to sensitivity of marginal utility to wealth, or a combination of both.

Existing research typically treats financial frictions and preference specifications as separate channels. On the financial side, credit-constraint models suggest that borrowing costs shape wealth accumulation, yet they rarely ask whether financial frictions are necessary for the emergence of threshold effects. By contrast, the consumption-savings literature highlights the role of utility curvature, yet without explicitly identifying structural breaks in wealth dynamics. Although reinforcement learning has recently been applied to consumption-savings problems [4], its potential for threshold identification remains largely unexplored. Consequently, the literature lacks a unified account to disentangle the behavioral imprint of preferences from the institutional footprint of financial markets.

The present work offers such an account by integrating deep reinforcement learning and structural break econometrics. A computational environment has been set up in this paper to find an optimal consumption-saving plan for a rational agent given wealth-dependent borrowing cost. At the same time, the interest-rate function is smooth and decreasing with wealth, without a threshold. A proximal policy optimisation (PPO) agent is employed to maximise log utility, and then the learned policy is used to generate a large-scale simulated panel of 5,000 independent trajectories over 1,000 periods. Hansen's panel threshold regression is used to detect whether the wealth-growth relationship shows a structural break endogenously [5]. Two more experiments are conducted to distinguish between behaviourally induced effects and institutional frictions. One excludes wealth-dependent borrowing costs, and the other replaces log utility with a CRRA function ($\gamma = 0.5$).

All the specifications have a statistically significant wealth threshold. It is mainly due to utility curvature; borrowing-cost asymmetry does not shift the estimated position of the threshold, and both the existence and location of the threshold are determined by non-linearity in marginal utility.

Rather than employing Reinforcement Learning for general-purpose numerical optimisation, the new system constructs optimal policy simulations and applies threshold econometrics to investigate why wealth changes at different times of accumulation. Therefore, people can distinguish between behavioral mechanisms and institutional frictions, and offer a computational approach that can be applied to other problems in AI and the social sciences.

The rest of this paper is as follows. Section 2 is in relation to the present studies. Section 3 presents the economic environment, the reinforcement learning framework and the econometric specification. Section 4 is the main experimental results, including the ablation study and utility comparison. Section 5 introduces the general implications and deficiencies, and Section 6 concludes. The technical content is in the Appendix.

2. Literature review

2.1. Wealth inequality and the Matthew Effect

The stylized fact that wealth is more concentrated than income has motivated a large body of theoretical and empirical work. Piketty famously argues that when the rate of return on capital exceeds the growth rate of the economy ($r > g$) [1], wealth inequality tends to widen mechanically.

More recent research has extended this discussion in different directions. Abatemarco, Lagomarsino and Spiganti develop a non-parametric empirical strategy to examine the relationship between the $r-g$ gap and income inequality in the United States [2], finding that this gap explains a substantial portion of the long-run trend. Jakurti offers a broader comparative perspective, demonstrating that the elasticity of top wealth shares with respect to the $r-g$ differential varies significantly across institutional contexts [3]. This suggests that the relationship is not purely mechanical but is instead shaped by policy and institutional environments. At the micro level, the

"Matthew Effect" certainly captures the intuition that the wealthy benefit from compounding advantages [6]. These studies explain why inequality persists, but they do not ask whether wealth accumulation changes discretely once households reach a particular level of wealth. The mechanism behind such a transition therefore remains unclear.

2.2. Financial frictions and wealth dynamics

Another line of research focuses on financial frictions. Ferlaino distinguishes between household-side frictions (borrowing constraints) and production-side imperfections [7], revealing that these distinct mechanisms drive consumption and wealth inequality through different channels. Building on a quantitative macroeconomic framework, Villalvazo demonstrates that the distribution of household leverage acts as a key transmission mechanism for financial shocks: during sudden stops [8], highly leveraged households suffer disproportionately larger wealth losses.

2.3. Reinforcement learning as a computational laboratory for economics

Reinforcement learning has increasingly been adopted to model forward-looking economic behavior. PPO provides a stable solution for dynamic optimization [9], while recent applications have shown that learned policies can reproduce realistic household decisions and macroeconomic interactions [5, 10]. Despite this progress, these studies focus primarily on policy learning rather than identifying structural changes in the resulting wealth dynamics.

2.4. Threshold regression and change-point detection

Hansen introduced panel threshold regression, which provides a data-driven approach for identifying endogenous structural breaks [5]. Subsequent work has extended the framework to more complex empirical settings [11, 12]. These developments provide the econometric foundation for examining nonlinearities in simulated wealth dynamics.

2.5. Positioning and novelty

The existing literature offers several explanations for wealth accumulation, yet it remains unclear whether the observed nonlinearity reflects preferences, financial frictions, or both. This paper addresses that question by combining reinforcement learning with threshold econometrics in a unified empirical framework.

3. Methodology

3.1. Economic environment

The paper constructs a discrete-time consumption-savings environment in which a single rational agent interacts with a stylized financial market. The agent's objective is to maximize the discounted sum of utilities from consumption over a finite horizon. The environment captures two core features, namely a trade-off between present consumption and future wealth accumulation and an effective return on savings that depends positively on current wealth through a smoothly varying interest rate function. The agent observes a two-dimensional state vector at each period $s_t = (W_t, t/T)$, where $W_t \geq 0$ denotes the current wealth and T is the maximum episode length.

3.1.1. Action space

The agent chooses a consumption ratio $a_t \in [0,1]$, which determines the fraction of wealth consumed in the current period. The remaining fraction $1 - a_t$ is saved and earns interest. This action space is continuous and one-dimensional.

3.1.2. Wealth-dependent interest rate

Savings receive a gross return $1 + r(W_t)$, where the interest rate function is:

$$r(W_t) = r_f + \kappa e^{-\lambda W_t} \quad (1)$$

Here $r_f \geq 0$ is a baseline risk-free rate, $\kappa > 0$ is the maximum additional spread charged to the poorest agents, and $\lambda > 0$ controls the speed at which the spread decays as wealth increases. The function is monotonically decreasing in wealth, satisfying

$$\lim_{W \rightarrow 0} r(W) = r_f + \kappa, \quad \lim_{W \rightarrow \infty} r(W) = r_f \quad (2)$$

This captures the idea that wealthier individuals face lower borrowing costs and have access to better investment opportunities. No discrete threshold is embedded in the interest rate schedule.

3.1.3. Labor income

Each period, the agent receives stochastic labor income y_t drawn from a log-normal distribution:

$$y_t \sim \text{LogNormal}(\mu_y, \sigma_y^2) \quad (3)$$

Income shocks are independent across time and agents.

3.1.4. Wealth transition

Given action a_t , the agent's wealth evolves as:

$$W_{t+1} = (1 - a_t) \cdot W_t \cdot (1 + r(W_t)) + y_t \quad (4)$$

Consumption is simply $C_t = a_t \cdot W_t$. An episode terminates after T periods or if wealth falls below a small positive threshold (bankruptcy).

3.1.5. Parameter calibration

The risk-free rate $r_f = 0.02$ is consistent with long-run U.S. financial data [10]. The income process parameters follow the buffer-stock saving literature [10]. The curvature parameters $\kappa = 0.30$ and $\lambda = 0.02$ produce a meaningful spread across the wealth distribution while maintaining a smooth, monotonic shape.

3.2. Reinforcement learning formulation

This study frames the agent's decision problem as a Markov Decision Process and solve it using Proximal Policy Optimization (PPO), a policy-gradient deep reinforcement learning algorithm introduced by [9].

3.2.1. Reward function

The agent receives a reward equal to the log utility of consumption:

$$R_t = \ln(C_t) = \ln(a_t \cdot W_t) \quad (5)$$

Log utility belongs to the CRRA family with coefficient of relative risk aversion equal to one, implying a strong precautionary saving motive. In the comparison experiment, the method is replacing log utility with CRRA preferences with $\gamma = 0.5$, i.e., $R_t = \frac{C_t^{1-\gamma}}{1-\gamma}$.

3.2.2. Objective

The agent selects a policy $\pi(a_t | s_t)$ to maximize the expected discounted sum of future rewards: $J(\pi) = \mathbb{E}_\pi \left[\sum_{t=0}^T \gamma^t R_t \right]$, where $\gamma \in (0, 1]$ is the discount factor.

3.2.3. Training procedure

This study employs a curriculum learning strategy. First, a PPO agent is pre-trained in a rich-start environment where initial wealth is drawn uniformly from [50, 100]. The agent learns the advantages of saving and the value of high-return states. Second, the trained policy is fine-tuned in the target environment where initial wealth is drawn from [5, 10]. This two-stage procedure stabilizes training and ensures performance across the entire wealth range.

3.2.4. Data generation

Once the policy is frozen, the paper generates a large-scale simulated panel: 5,000 independent episodes, each with a distinct random seed determining initial wealth and income shocks, for $T = 1000$ periods. For each episode i and period t , the paper records wealth $W_{i,t}$ and consumption $C_{i,t}$. The growth rate is:

$$g_{i,t} = \ln \left(\frac{W_{i,t+1}}{W_{i,t}} \right) \quad (6)$$

The resulting dataset contains approximately 2.5 million observations after removing outliers and episodes terminated by bankruptcy.

3.2.5. Nature of the panel

This is a trajectory ensemble from a single optimal policy; individual heterogeneity arises solely from differences in initial endowments $W_{i,0}$ and the accumulation of idiosyncratic income shocks. Fixed effects μ_i in the subsequent panel regressions therefore absorb the level effect of initial

conditions rather than representing structural differences in preferences or constraints across agents. The threshold the paper identifies captures a state-dependent break under a common decision rule, not one driven by unobserved agent-level heterogeneity.

3.3. Panel threshold regression

To test for a structural break in the wealth–growth relationship, the paper employs the panel threshold regression framework of Hansen [5]. The baseline specification is:

$$g_{i,t} = \beta_1 \cdot \log W_{i,t} \cdot 1\{\log W_{i,t} \leq \gamma\} + \beta_2 \cdot \log W_{i,t} \cdot 1\{\log W_{i,t} > \gamma\} + \mu_i + \epsilon_{i,t} \quad (7)$$

where γ is the threshold parameter, $1\cdot$ is the indicator function, and μ_i captures agent-specific fixed effects. β_1 measures the marginal effect of log wealth on the growth rate below the threshold; β_2 measures the effect above.

3.3.1. Estimation

For each candidate threshold γ , the paper estimates by OLS after the within transformation to eliminate μ_i . The optimal threshold $\hat{\gamma}$ minimizes the concentrated RSS:

$$\hat{\gamma} = \arg \min_{\gamma} RSS(\gamma) \quad (8)$$

This study searches over a grid covering the 10th to 99th percentiles of the log-wealth distribution to ensure that the global minimum is captured.

3.3.2. Inference

The paper tests $H_0 : \beta_1 = \beta_2$ using an F-statistic constructed from the RSS of the linear and threshold models. Under the null, this statistic follows an approximate F-distribution.

3.3.3. Ablation and comparison

To decompose the sources of the threshold, this paper run identical specifications on two alternative environments: one is a flat-cost environment where $r(W) \equiv r_f + \kappa e^{-\lambda \cdot 50}$; another one is the asymmetric-cost environment with CRRA utility ($\gamma = 0.5$). Comparing thresholds isolates preference curvature from financial asymmetry.

4. Results

4.1. Descriptive statistics

Figure 1 Wealth growth rate plots the mean wealth growth rate by wealth decile in the baseline simulation: the poorest decile exhibits an average growth rate of +0.845, while the richest decile shows −0.849. The downward-sloping pattern is consistent with optimal consumption smoothing under log utility. No visible discontinuity appears in Figure 1, suggesting that any structural break is more likely to involve a change in slope than a jump in the growth rate.




Figure 1. Wealth growth rate (picture credit: original)

4.2. Baseline threshold regression

Hansen's threshold regression identifies a structural break at log wealth $\hat{\gamma} = 2.1003$ [5], corresponding to a wealth level of 8.17 (Figure 2 RSS profile for baseline regression).

Below the threshold, each one-unit increase in log wealth is associated with a steeper decline in growth.; once the threshold is crossed, this dampening effect moderates markedly. The F-statistic reaches 213.5 ($p < 0.001$), rejecting the null hypothesis of no threshold. The corresponding 95% confidence interval for the threshold, constructed via Hansen's likelihood-ratio procedure, spans [7.60, 8.69] in raw wealth terms (equivalent to [2.028, 2.162] in log-wealth units) [5].




Figure 2. RSS profile for baseline regression (picture credit: original)

4.3. Mechanism decomposition

The ablation exercise with flat borrowing costs yields a threshold of 7.78, with an F-statistic of 208.6 and $p < 0.001$. The CRRA comparison with $\gamma = 0.5$ produces a threshold of 7.96, with an F-statistic of 220.2 and $p < 0.001$. All three estimates cluster tightly within the 7.8–8.2 range and remain

statistically indistinguishable. The 95% confidence interval for the baseline threshold spans 7.60 to 8.69. Both the ablation estimate of 7.78 and the CRRA estimate of 7.96 fall well within this interval, confirming that the three values do not differ substantively.

The three threshold estimates are remarkably similar despite differences in the underlying environments. Equalizing borrowing costs leaves the estimated threshold essentially unchanged, whereas replacing the utility specification also produces only a modest shift. Taken together, these results are consistent with the view that the observed threshold is primarily associated with preference curvature rather than borrowing-cost asymmetry.

5. Discussion

The central empirical finding is that a statistically significant structural break exists in the wealth–growth relationship, but its origin is not what the intuitive narrative of "wealth begets wealth" would suggest. The break is primarily driven by utility curvature, while borrowing-cost asymmetry does not materially affect its location.

5.1. Interpretation of the mechanism

Under log utility, agents optimally shift from accumulation to steady-state consumption, yielding a downward-sloping wealth–growth profile—an implication of consumption smoothing under diminishing marginal utility [10]. This downward trend, and the wealth level at which the transition occurs, are unaffected by financial frictions. The structural break thus reflects a moderation of the wealth–growth gradient rooted in preference curvature, not in borrowing-cost asymmetry.

5.2. Relation to the $r > g$ hypothesis

The paper's analysis contributes a micro-foundation: individual wealth–growth profiles exhibit breaks even without systematic return differentials. The threshold arises from the nonlinearity of marginal utility, not from financial frictions. This suggests that the "wealth begets wealth" phenomenon may be more deeply rooted in preference structures than in institutional asymmetries. The results complement the aggregate-level findings of Abatemarco et al. and Jakurti by showing that [2, 3], at the micro level, the relationship between wealth and its own growth rate is non-linear and regime-dependent even in the absence of the $r > g$ mechanism.

5.3. Policy implications

The wealth threshold remains stable after removing borrowing-cost asymmetry. Thus, policies like microfinance or rate caps can improve welfare without moving the transition wealth. To shift the threshold, preference curvature must change—yet this is only theoretical. Villalvazo finds that macroprudential tools affect volatility, but not the breakpoint [8].

5.4. Limitations and future work

The paper's panel is a single-policy ensemble, whereas real households have heterogeneous preferences and constraints, so future work could extend the RL framework to include heterogeneous agents by building on the multi-agent RL architectures explored by Brusatin et al. and Garrido-Merchán et al. [13, 14]. Moreover, the current environment is partial equilibrium, and only two utility specifications are tested. In addition, validation against real panel data (e.g., CHFS,

PSID) remains future work, though the random change point framework of Li, Gao, Lv and Tong [12], applied to household survey data, offers a promising direction for empirical validation. Accordingly, future extensions should introduce heterogeneous agents, general equilibrium feedback, and empirical validation with household panel data.

6. Conclusion

This paper identifies a statistically significant structural break in the wealth–growth relationship at wealth 8.17 (95% CI: [7.60, 8.69]). Through ablation and utility comparison, it shows that this threshold is invariant to borrowing-cost asymmetry and utility specification, with all three experimental thresholds falling within the same confidence interval. The threshold is driven by the curvature of the utility function—the nonlinearity of marginal utility that causes optimal consumption–savings behavior to shift from accumulation to consumption as wealth grows. Financial frictions do not materially affect the estimated threshold in the present environment. The paper demonstrates how reinforcement learning with panel threshold analysis can rigorously test intuitive hypotheses about the micro-foundations of wealth inequality. Future work should extend this framework to heterogeneous agents with diverse preferences, general equilibrium settings with endogenous interest rates, and empirical validation using real-world household panel data.

References

- [1] Piketty, T. (2014). *Capital in the twenty-first century*. Harvard University Press. Cambridge, MA.
- [2] Abatemarco, A., Lagomarsino, E., & Spiganti, A. (2025). Rising income inequality and the impact of $r-g$ in the United States. *Review of Income and Wealth*, 71(1), e70000.
- [3] Jakurti, E. (2025). A tale of two rates: Return on capital, economic growth, and wealth concentration in the long run. *Review of Political Economy*. <https://doi.org/10.1080/09538259.2025.2492268>
- [4] Kaplowitz, B. (2025). Reinforcement learning and consumption-savings behavior. *arXiv*, arXiv:2510.20748.
- [5] Hansen, B. E. (1999). Threshold effects in non-dynamic panels: Estimation, testing, and inference. *Journal of Econometrics*, 93(2), 345–368.
- [6] Merton, R. K. (1968). The Matthew effect in science. *Science*, 159(3810), 56–63.
- [7] Ferlaino, F. (2025). Effects of different financial frictions on households. *Working Papers in Public Economics No. 263*, Sapienza University of Rome.
- [8] Villalvazo, S. (2026). Inequality and asset prices during sudden stops. *Journal of Monetary Economics*, 157, 103872.
- [9] Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv*, arXiv:1707.06347.
- [10] Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly Journal of Economics*, 112(1), 1–55.
- [11] Li, J., & Yan, H. (2024). Uniform inference in high-dimensional threshold regression models. *arXiv*, arXiv:2404.08105.
- [12] Li, M., Gao, L., Lv, G., & Tong, X. (2024). Random change point model with an application to the China Household Finance Survey. *Science China Mathematics*, 67(6), 2373–2386.
- [13] Brusatin, S., Padoan, T., Coletta, A., Delli Gatti, D., & Glielmo, A. (2024). Simulating the economic impact of rationality through reinforcement learning and agent-based modelling. *arXiv*, arXiv:2405.02161.
- [14] Garrido-Merchán, E. C., López-López, Á., & Coronado-Vaca, M. (2024). Deep reinforcement learning agents for strategic production policies in microeconomic market simulations. *arXiv*, arXiv:2410.20550.