Fiscal Stabilisers, Minsky Dynamics, and Distributional Outcomes in a Keynesian Agent-Based Model

Baseline Simulation Evidence from a 10-Year Closed-Economy Run

Evint Leovonzko*

Abstract

We develop, calibrate, and analyse a daily-frequency, single-sector Keynesian agent-based model (ABM) of a closed economy populated by 10,000 households, 1,000 firms, and 10 commercial banks operating alongside a fiscal authority and a Taylor-rule central bank. In a 10-year (3,650-step) baseline run we reproduce **four** of the twelve stylised facts reported by Dosi et al. [13]: a nominal GDP compound annual growth rate of 4.6 %, a mean year-on-year growth rate of 4.6 %, negatively skewed annual growth fluctuations ($\hat{\gamma}_1 = -0.59$), and a wage share of GDP of 72.8 %. The model is deliberately restricted to a single consumption-good sector and is therefore not designed to reproduce capital-market or multi-sector phenomena. The economy reaches the zero lower bound (ZLB) from day 1 and remains there for a median 93.7 % of simulation periods across 50 Halton-seed Monte Carlo draws (95 % CI: [6.0 %, 96.1 %]), mirroring the post-1998 Japanese and post-2013 Euro-area experience—a novel result in the K+S ABM literature. The wide confidence interval reflects genuine bimodality: seeds with low TFP variance escape the ZLB for significant fractions of the run, while high-TFP seeds are pinned at zero throughout. GDP expands at a steady 4.6 %/yr under the recalibrated R&D specification; the earlier parameterisation that produced a Minsky-reversal bankruptcy spike is superseded. We document five structural departures from the empirical target—near-zero unemployment, supply-side CPI deflation, suppressed business-cycle volatility and persistence, compressed wealth inequality, and persistent ZLB—and trace each to identifiable modelling choices in the single-sector design. Two of these—suppressed growth volatility (0.12 % vs. empirical 2–5 %) and nega-

*Corresponding author
Email address: evint.koo@gmail.com (Evint Leovonzko)

tive growth autocorrelation ($\hat{\rho}_1 = -0.18$ vs. empirical $\hat{\rho}_1 > 0.70$)— arise directly from the smooth, near-deterministic TFP growth path produced by the recalibrated R&D parameters, and establish the boundary conditions for when endogenous cycles re-emerge as the model is extended toward the full two-sector design.
Keywords: Agent-based model, Keynesian macroeconomics, Business cycles, Zero lower bound, Minsky dynamics, Firm bankruptcy, Fiscal policy
JEL: C63, E12, E32, E44, E62, G01

---

1. Introduction

The 2007–09 Global Financial Crisis and the subsequent decade of secular stagnation demonstrated the inadequacy of models in which the macroeconomy converges rapidly to a unique rational-expectations equilibrium. In those frameworks, large and persistent output contractions require correspondingly large exogenous shocks—a feature Caballero [7] memorably described as the "pretense-of-knowledge syndrome." By contrast, agent-based computational models of the macroeconomy (macro-ABMs) generate large-amplitude fluctuations endogenously from the local interactions of heterogeneous, boundedly rational agents operating in decentralised markets [18]. The macro-ABM literature has produced several benchmark platforms alongside the K+S programme, including the Eurace@Unibi model [11] and the Agent-Based Stock-Flow-Consistent framework of Caiani et al. [8]; together these establish a common set of validation targets that any new macro-ABM should engage with.

The Keynes + Schumpeter (K+S) programme initiated by Dosi et al. [15] and extended in Dosi et al. [13, 14] has established a benchmark for this approach. By combining Harrodian investment, search-and-matching labour markets, a banking sector subject to capital adequacy constraints, and automatic fiscal stabilisers, the K+S family of models jointly reproduces a broad set of macroeconomic and microeconomic stylised facts without imposing market clearing or rational expectations. Crucially, these properties emerge for a wide range of parameter values, suggesting that they are structural features of the Keynesian coordination mechanism rather than artefacts of fine-tuning [14].

The present paper makes three contributions to this literature.

First,. this paper constructs a *minimal single-sector null model* designed to isolate by subtraction which stylised facts require two-sector structure. The design deliberately omits the capital-good sector, equity market, and nominal price anchor of the full K+S architecture. Implemented in Rust, a 10-year, 10,000-agent run completes in under 30 seconds, enabling high-throughput Monte Carlo experiments that are computationally prohibitive in larger K+S implementations. The eight empirical misses documented in Section 5 are therefore the central finding: each identifies a structural feature the single-sector design cannot replicate, providing a precise diagnostic for what the two-sector extension must supply.

Second,. this paper establishes a systematic mapping between structural design choices and empirical misses. Near-zero unemployment—a direct consequence of the missing capital-income channel and the absence of a statutory wage floor—is the proximate cause of four of the five structural departures (Proposition 3). The two business-cycle departures from the empirical record (growth volatility and autocorrelation) additionally trace to the recalibrated R&D parameters, which produce near-deterministic TFP growth. This hierarchy converts a catalogue of failures into a single diagnostic: reproducing business-cycle dynamics requires either restoring volatile TFP growth or introducing the two-sector capital-good pricing that provides a nominal anchor.

Third,. the most striking result of the baseline simulation is that the zero lower bound binds in the *large majority* of simulation periods. Across 50 Halton-seed Monte Carlo draws the ZLB frequency reaches a median of 93.7 % (95 % CI: [6.0 %, 96.1 %]; Figure 3), with genuine bimodality reflecting a TFP-variance threshold effect (Section 6.5). This result is novel in the K+S ABM literature and mirrors the post-1998 Japanese and post-2013 Euro-area experience [28, 16]. It provides a sharp, testable distinction from the K+S two-sector model: capital-good sector pricing acts as a nominal anchor that makes ZLB episodes occasional rather than chronic [14]; its absence in the single-sector design makes the ZLB the default state.

Organisation.. Section 2 describes the agents and their behavioural rules. Section 3 presents the calibration and shock schedule. Section 4 reports the main simulation results. Section 5 compares simulated moments with the Dosi et al. [13] benchmarks. Section 6 discusses structural mechanisms and calibration pathways. Section 7 concludes.

Unless otherwise stated, moments are 50-seed Monte Carlo means with 95 % CI from Halton-sequence seeds (Figure C.11). The seed-42 single-run results are the primary reference for table entries; MC statistics are reported where they provide useful distributional context.

2. The Model

The economy is a single-sector closed system. Time is discrete at the *daily* frequency; the simulation horizon is $T = 3{,}650$ days (10 calendar years).[1] A single consumption good is produced by firms, allocated through decentralised goods and labour markets, and consumed by households or purchased by the government. There is no capital-good sector and no international trade.

2.1. Households

Households follow a buffer-stock consumption rule [9], holding precautionary savings as a buffer against income uncertainty, and supplying labour inelastically subject to a reservation-wage friction.

There are $N_H = 10{,}000$ households indexed $i \in \mathcal{H}$. Each household holds *cash* $m_{i,t}$, bank *savings* $s_{i,t}$, and may carry bank *debt* $d_{i,t}$. Net wealth is $w_{i,t} = m_{i,t} + s_{i,t} - d_{i,t}$.

[1]All behavioural parameters are expressed in *annual* units in the configuration file and converted to daily equivalents at run time (e.g. an annual interest rate of r becomes a daily accrual of $r/365$). This convention eliminates unit mismatches common in daily-frequency ABMs.

Income

Household i's daily effective income is

$$y_{i,t} = \underbrace{\tilde{\omega}_{i,t}}_{\text{labour income}} + \underbrace{\tau_t}_{\text{social transfer}}, \tag{1}$$

where $\tilde{\omega}_{i,t}$ is the daily wage for employed workers and the daily unemployment benefit b_t for the unemployed:

$$\tilde{\omega}_{i,t} = \begin{cases} W_{j(i),t}/365 & \text{if employed at firm } j(i), \\ b_t = \beta \bar{W}_t/365 & \text{if unemployed.} \end{cases} \tag{2}$$

Here $W_{j(i),t}$ is the annualised offered wage of the employing firm, β is the replacement rate, $\bar{W}_t$ is the economy-wide average annualised wage, and τ_t is a daily social transfer set by the government (Section 2.4). Both b_t and τ_t are paid *before* the daily goods market opens, ensuring that transfer income circulates through the consumption channel.

Consumption budget

Disposable income after the daily fraction of income tax θ^H is $\tilde{y}_{i,t} = y_{i,t}(1 - \theta^H/365)$. The *target* consumption budget is

$$c^*_{i,t} = \tilde{y}_{i,t}\Big[1 - \sigma(1+\rho)\Big], \tag{3}$$

where $\sigma \in (0,1)$ is the savings propensity and $\rho > 0$ the precautionary risk-aversion coefficient. The effective marginal propensity to consume is therefore $\text{MPC} = 1 - \sigma(1+\rho)$. Habit persistence enters through a partial-adjustment rule:

$$c_{i,t} = \eta\, c_{i,t-1} + (1-\eta)\, c^*_{i,t}, \tag{4}$$

where $\eta \in [0,1]$ is the habit coefficient. To avoid an artificial demand ramp-up in the first 30 days, $c_{i,0}$ is pre-warmed to the steady-state value $c^*_{i,0}$ at initialisation. The realised budget is bounded: $c_{i,t} \in [0, m_{i,t}]$.

Labour supply

Each unemployed household holds an annualised *reservation wage* $W^r_{i,t}$. On each day the household samples k^L vacancies drawn uniformly without replacement from the current vacancy pool and accepts the highest posted wage $W_{j,t} \geq W^r_{i,t}$, provided a Bernoulli draw with success probability $1-\xi$ succeeds (matching friction). If no match occurs the reservation wage decays as

$$W^r_{i,t+1} = W^r_{i,t}\left(1 - \frac{\delta_W}{365}\right), \tag{5}$$

where δ_W is the annual decay rate [33]. Upon accepting an offer the reservation wage is reset to the accepted wage.

2.2. Firms

Firms produce a single consumption good using a Cobb–Douglas technology, set prices via an adaptive markup rule [17], hire and fire workers in response to demand signals, and finance investment from retained profits supplemented by bank credit.

Initially $N_F = 1{,}000$ firms are active. Firm j is characterised by its total factor productivity $A_{j,t}$, physical capital $K_{j,t}$, current employment $L_{j,t}$, annualised wage offer $W_{j,t}$, output price $P_{j,t}$, and markup $\mu_{j,t}$.

Production

Output follows a Cobb–Douglas technology:

$$Y_{j,t} = A_{j,t}\, L^{\alpha}_{j,t}\, K^{\beta}_{j,t}, \qquad \alpha + \beta = 1. \tag{6}$$

The Cobb–Douglas specification with $\alpha = 0.60$, $\beta = 0.40$ is consistent with the empirical labour-income shares documented for OECD economies [13, 35]. Capital evolves as

$$K_{j,t+1} = K_{j,t}\left(1 - \frac{\delta_K}{365}\right) + I_{j,t}, \tag{7}$$

where δ_K is the annual depreciation rate and investment $I_{j,t} = \chi \max(\pi_{j,t-1}, 0)$ is a fixed fraction χ of last period's non-negative profit. R&D spending equals $\phi \cdot \text{Rev}_{j,t}$; with probability $\zeta\phi$ this raises $\ln A_{j,t}$ by a draw from $\mathcal{N}(0, s_A^2)$.

Pricing

Prices are set via a markup over average unit cost:

$$P_{j,t} = (1 + \mu_{j,t}) \frac{W_{j,t}}{A_{j,t}}. \tag{8}$$

The markup adjusts adaptively at daily rate δ_μ:

$$\mu_{j,t+1} = \begin{cases} \mu_{j,t} + \delta_\mu & \text{if demand/capacity} > \theta_f^H, \\ \max(\mu_{j,t} - \delta_\mu,\, 0) & \text{otherwise.} \end{cases} \tag{9}$$

Proposition 1 (Balanced-growth deflation)**.** *Along a balanced-growth path in which firm-level TFP grows at mean rate g_A and the economy-wide nominal wage grows at rate g_W, the aggregate price level falls at rate*

$$g_P \;\approx\; g_W - g_A. \tag{10}$$

Proof sketch. From the markup pricing rule (8), $P_{j,t} = (1 + \mu_{j,t})W_{j,t}/A_{j,t}$. Taking logs and first differences, and noting that the markup $\mu_{j,t}$ is stationary around its target in a balanced growth path, gives $\Delta \ln P_{j,t} \approx \Delta \ln W_{j,t} - \Delta \ln A_{j,t}$. Aggregating across firms yields $g_P \approx g_W - g_A$. In the baseline calibration, $g_A \approx 12\,\%$/yr (implied by R&D parameters $\phi = 0.02$, $\zeta = 0.10$) and $g_W \approx 2\,\%$/yr (wage drift at low utilisation), so $g_P \approx -10\,\%$/yr, consistent with the simulated deflation in Table 8. □

Hiring and wage setting

Vacancies are posted whenever the demand-to-capacity ratio exceeds θ_f^H; workers are dismissed when it falls below the firing threshold θ_f^F. Annualised wage offers adjust as

$$W_{j,t+1} = W_{j,t} + \delta_W^f \, \bar{W}_t \cdot \mathbf{1}_{[\text{tight}]} - \delta_W^f \, \bar{W}_t \cdot \mathbf{1}_{[\text{slack}]} \, (1 - \phi_W), \tag{11}$$

where δ_W^f is the daily drift rate, $\bar{W}_t$ is the economy-wide average wage, and $\phi_W \in [0, 1]$ captures downward nominal rigidity. Equation (11) is estimated in natural-rate units: the drift $\delta_W^f \bar{W}_t$ is a daily increment proportional to the level of wages, not to the firm's own wage, preventing explosive relative-wage dynamics.

Profits and bankruptcy

Daily gross profit is $\pi_{j,t} = P_{j,t} Q_{j,t} - (W_{j,t}/365) L_{j,t} - r_{j,t}^B D_{j,t} - \theta^C \max(\pi_{j,t}^{\text{pre-tax}}, 0)$, where $Q_{j,t}$ is quantity sold, $r_{j,t}^B$ the loan rate, $D_{j,t}$ outstanding debt, and θ^C the corporate tax rate. Firm j is declared bankrupt if its equity ratio $e_{j,t} = (\text{assets} - \text{liabilities})/\text{assets}$ falls below the threshold $e^* = -0.5$.

Entry

A Poisson number of entrants with intensity

$$\lambda_t = \nu \, N_F^{\text{alive}}(t) \tag{12}$$

replaces exiting firms, where ν is the daily entry rate per surviving firm.[2] Entrant TFP is drawn from $\text{LogNormal}(0, s_A^2)$; the initial wage offer is set to the marginal product of labour implied by a single employee and average firm capital.

[2]Using N_F^{alive} rather than N_F^{total} (which includes deceased-but-stored firm objects) prevents exponential expansion of the firm vector—a subtle but critical implementation detail.

2.3. Commercial Banks

Banks intermediate between household depositors and firm borrowers, subject to a Basel II minimum capital adequacy ratio. Their fragility is the transmission channel for the Minsky cycle: when bank equity is destroyed, credit supply contracts, propagating firm distress into the real economy.

Ten commercial banks accept household deposits, extend corporate loans, and are subject to a minimum capital adequacy ratio (CAR). Firm j borrows from a randomly sampled bank at daily interest rate

$$r^B_{j,t}/365 = r_t/365 + \frac{\zeta^B}{365}\max(0, -e_{j,t}), \tag{13}$$

where r_t is the central bank policy rate and ζ^B is a risk-premium slope on negative firm equity. When a bank's own equity ratio falls below the CAR, new lending is suspended and the government recapitalises the bank at cost.

2.4. Government

The government operates automatic fiscal stabilisers: unemployment benefits, social transfers, and bank recapitalisations that expand counter-cyclically, cushioning demand in downturns without requiring discretionary intervention.

The government collects income tax at rate θ^H from households and corporate tax at rate θ^C from firms. Daily expenditure is

$$G_t = n^u_t\, b_t + N_H\, \tau_t + \mathcal{B}_t, \tag{14}$$

where n^u_t is the count of unemployed workers, $\mathcal{B}_t$ is bank-bailout expenditure, and the social transfer is set by the automatic stabiliser rule

$$\tau_t = \frac{\psi\, \bar{W}_t}{365}\Big[1 + \phi_G(u_t - u^*)\Big], \tag{15}$$

with target unemployment rate $u^* = 0.05$, stabiliser strength ϕ_G, and transfer generosity ψ.

Fiscal deficits accumulate as government debt financed by bond issuance; a soft ceiling of $\bar{d} = 3\,\%$ of GDP is imposed but not enforced when the automatic stabiliser is triggered.

2.5. Central Bank

The central bank sets the overnight rate according to a standard Taylor rule [39], subject to a zero lower bound. Quantitative easing is available at the ZLB but, as demonstrated in Section 6.5, does not alter the real equilibrium in this single-sector design.

The policy rate follows a standard Taylor rule:

$$r_t = \max(0,\ \bar{r} + \psi_\pi(\pi_t - \pi^*) + \psi_y\, \tilde{y}_t)\,, \tag{16}$$

where $\bar{r} = 0.02$ is the neutral rate, π_t is the 30-day annualised inflation rate, $\pi^* = 0.02$ the inflation target, and $\tilde{y}_t$ the output gap estimated from a rolling HP-filtered potential GDP series. The floor in (16) imposes the zero lower bound. The ceiling is $r^{\max} = 0.15$. Quantitative easing is available when $r_t = 0$: the central bank purchases assets at rate q per period, expanding bank reserves.

2.6. Markets

Goods market.. Each day households draw a sample of k^G firms uniformly at random and purchase from the cheapest supplier until their budget is exhausted or supply is rationed. The aggregate demand equals total household consumption budgets plus government spending. The daily unmet-demand fraction υ_t is recorded as an indicator of supply-side rationing.

Labour market.. The labour market clears through a two-phase algorithm: (i) firms post vacancies with associated wage offers; (ii) unemployed households are shuffled randomly and each draws k^L vacancies, accepting the best offer above its reservation wage conditional on a match succeeding with probability $1 - \xi$. Unmatched workers receive benefits and decay their reservation wage.

2.7. Theoretical Foundations of Agent Decision Rules

The bounded-rational rules described above can each be related to a theoretically optimal counterpart from the macroeconomics literature. This section states these foundations; the companion *optimal* agent variant activates them in simulation and is compared with the baseline in Section 4.12.

Household consumption (Euler equation).. Hall [24] shows that under CRRA utility, optimal consumption satisfies

$$c_t^* = \Big(\beta(1+r)\,\hat{E}[c^{-\gamma}]\Big)^{-1/\gamma}, \tag{17}$$

where β is the discount factor, r the deposit rate, and γ the coefficient of relative risk aversion. The bounded-rational rule replaces this with a habit-income mixture.

Reservation wage (McCall search).. McCall [31] derives the reservation wage W^* as the fixed point of

$$W^* = b + \frac{\beta}{1+\delta_W}\,E[\max(w - W^*,\,0)], \tag{18}$$

where b is the unemployment benefit and δ_W the offer-arrival hazard. The bounded-rational rule decays W^* geometrically.

Firm investment (Tobin's q).. Hayashi [25] establishes the q-theory investment rule $I/K = (q-1)/\phi_{\text{adj}}$, where $q = V/K$ is the ratio of firm value to replacement cost and ϕ_{adj} an adjustment cost. Under uncertainty, Dixit and Pindyck [12] show that the investment threshold rises to $q^* = \beta_1/(\beta_1 - 1) > 1$, where β_1 is the positive root of the value-matching equation.

Firm pricing (Lerner markup).. Lerner [30] derives the optimal markup as $\mu^* = 1/(\varepsilon - 1)$, where ε is the absolute price elasticity of demand estimated from rolling OLS of $\Delta \ln p$ on $\Delta \ln q$. The price-flexibility bound κ_P limits the day-to-day change in unit costs passed through to prices.

Firm leverage (Kalecki principle).. Kalecki [26] argues that investment is bounded by the principle of increasing risk: $I \leq K \cdot (\text{equity/assets})^{\phi_K}$. The bounded-rational rule uses a fixed investment rate χ.

Bank lending (Stiglitz–Weiss and Bernanke–Gertler).. Stiglitz and Weiss [38] show that, with adverse selection, the optimal loan rate is $r^* = r_{\text{safe}} + 1/\hat{\lambda}$, where $\hat{\lambda}$ is the MLE hazard from observed default outcomes. Bernanke and Gertler [5] add a convex leverage premium, and Kiyotaki and Moore [27] impose an asset-price collateral ceiling.

Fiscal policy (Barro tax smoothing).. Barro [4] derives the optimal tax rate as the permanent-income share: $\tau^* = \bar{G}/\bar{Y}$, where $\bar{G}$ and $\bar{Y}$ are the long-run averages of government spending and output.

Monetary policy (Woodford and Eggertsson–Woodford).. Woodford [40] derives the welfare-loss-minimising Taylor rule with inflation coefficient $\psi_\pi^* = 1 + \lambda_y/(\lambda_\pi \kappa^2)$. Eggertsson and Woodford [16] show that at the zero lower bound the optimal commitment policy maintains accommodation for $T^* = \ln(\pi^*/\pi)/[\kappa(\psi_\pi - 1)]$ periods.

3. Calibration

Parameter values are reported in Tables 1–5. The calibration follows a three-tier strategy: *(i)* parameters shared with the K+S two-sector model are taken directly from Dosi et al. [13]; *(ii)* standard macroeconomic quantities (tax rates, depreciation, Basel CAR) are set to their conventional empirical values; and *(iii)* parameters without a K+S or empirical anchor (entry rate, investment rate, matching friction) are set to produce plausible initial conditions and are flagged as free parameters subject to sensitivity analysis (Appendix Appendix D).[3]

[3] The sensitivity appendix varies five free parameters—matching friction ξ, hiring threshold θ_f^H, reservation-wage decay δ_W, investment rate χ, and price-flexibility bound κ_P—over a four-point grid with five Halton seeds per grid point (100 runs per parameter). Each free parameter is benchmarked as follows: entry rate $\nu = 0.0002$/day implies a gross firm entry rate of $\approx 7\,\%$/yr, consistent with US Longitudinal Business Database estimates [**?**]; investment rate $\chi = 0.10$ is calibrated to match a pre-crisis investment-to-profit ratio within the range documented by Dosi et al. [13]; matching friction $\xi = 0.10$ targets a vacancy-filling rate of 90 %, toward the upper end of the **?**] estimates for the US.

A limitation of the one-at-a-time (OAT) sensitivity design is that it cannot identify the *joint* parameter manifold consistent with any given target moment. The free parameters interact: reducing ξ and increasing θ_f^H both raise vacancy-filling rates, but their joint effect on the unemployment rate is non-additive. In particular, the OAT ξ-sweep reported in Section 4.9 confirms that raising ξ from 0.10 to 0.60 leaves unemployment at approximately 0.006 %—not because the parameter is inconsequential per se, but rather because the pool of job-seekers is so small that a high rejection rate cannot generate a meaningful unemployment stock. A joint grid search over $(\xi, \theta_f^H, \delta_W)$ is required to identify the parameter combination that simultaneously delivers positive unemployment and a realistic Beveridge-curve relationship. This three-dimensional grid search is deferred to future work; the OAT sweeps reported here establish the boundaries of each parameter's individual effect and confirm that the baseline is in the supply-constrained rather than search-friction-constrained regime.

The household wealth distribution at $t = 0$ is drawn from a Pareto tail with index $\alpha_H = 1.5$ (consistent with empirical distributions in Gabaix 20); firm sizes are drawn from a Pareto with $\alpha_F = 1.2$ [2, 36]. The total money supply at initialisation is $1,000,000, split 50 %/30 %/20 % across households, firms, and banks. The baseline consumption budget pre-warming sets $c_{i,0} = \tilde{y}_{i,0} \cdot \text{MPC}$ for all households to suppress the mechanical demand ramp-up from habit persistence in the first 30 days.

Table 1: Simulation parameters

Parameter	Symbol	Value
Households	N_H	10,000
Firms (initial)	N_F^0	1,000
Banks	N_B	10
Simulation days	T	3,650
Random seed	—	42
CSV flush interval (days)	—	30
Expectation window (days)	κ	30

Table 2: Household and labour market parameters

Parameter	Symbol	Value	Reference
Risk aversion	ρ	0.20	—
Savings propensity	σ	0.05	MPC = 0.94
Consumption habit	η	0.70	Fuhrer [19]
Goods sample size	k^G	8	Dosi et al. [13]
Labour search sample	k^L	5	Dosi et al. [13]
Matching friction	ξ	0.10	—
Reservation wage decay (ann.)	δ_W	0.05	Dosi et al. [13]
Debt tolerance	—	2.0	—
Wealth Pareto tail	α_H	1.50	Gabaix [20]

Table 3: Firm parameters

Parameter	Symbol	Value	Reference
Initial markup	μ_0	0.20	Dosi et al. [13]
Markup step	δ_μ	0.02	Dosi et al. [13]
Hiring threshold	θ_f^H	0.90	Dosi et al. [13]
Firing threshold	θ_f^F	0.50	Dosi et al. [13]
Investment rate	χ	0.10	Dosi et al. [13]
Capital depreciation (ann.)	δ_K	0.05	OECD avg. 5 %/yr
Bankruptcy threshold	e^*	−0.50	Greenwald and Stiglitz [21]
Dividend payout ratio	—	0.40	free parameter
R&D spending rate	ϕ	0.02	Dosi et al. [13]
R&D efficiency	ζ	0.10	—
Entry rate (daily, per alive firm)	ν	0.0002	≈ 7 %/yr
Cobb–Douglas labour share	α	0.60	—
Cobb–Douglas capital share	β	0.40	—
Wage drift rate (daily)	δ_W^f	0.0002	≈ 7 %/yr max
Wage rigidity	ϕ_W	0.40	Dosi et al. [13]
Productivity dispersion (s.d.)	s_A	0.30	Dosi et al. [13]
Firm size Pareto tail	α_F	1.20	Axtell [2]

Table 4: Banking sector parameters

Parameter	Symbol	Value	Reference
Capital adequacy ratio (CAR)	—	0.08	Basel II
Lending appetite	—	0.80	—
Risk-premium slope	ζ^B	0.50	—
Max loan-to-equity ratio	—	0.70	—
Deposit spread	—	0.01	—
Bank sample (firms per bank)	k^B	3	—

Table 5: Fiscal and monetary policy parameters

Parameter	Symbol	Value	Reference
Income tax rate (annual)	θ^H	0.25	—
Corporate tax rate (annual)	θ^C	0.21	—
Unemployment replacement rate	β	0.40	Dosi et al. [13]
Transfer generosity	ψ	0.50	—
Stabiliser strength	ϕ_G	0.30	—
Target unemployment rate	u^*	0.05	5 %
Deficit ceiling (GDP share)	$\bar{d}$	0.03	SGP
Taylor weight on inflation	ψ_π	1.50	Taylor [39]
Taylor weight on output gap	ψ_y	0.50	Taylor [39]
Neutral interest rate (ann.)	$\bar{r}$	0.02	—
Inflation target (ann.)	π^*	0.02	—
ZLB floor	—	0.00	—
Interest rate ceiling (ann.)	—	0.15	—
QE asset-purchase rate	q	0.01	—

Exogenous shocks.. To test the model's propagation mechanism, two aggregate shocks are imposed (Table 6). These are sized to be large enough to leave a clear signal in the aggregate time series while remaining within the range considered in Dosi et al. [13]. Together, they allow the paper to assess the model's propagation mechanisms—automatic-stabiliser cushioning after the demand shock, and Minsky-reversal credit dynamics after the banking crisis—while preserving comparability with the K+S validation benchmark.

Table 6: Exogenous shock schedule

Day	Year	Type	Magnitude
500	1.37	Demand shock (consumption scale-back)	−30 %
1200	3.29	Banking crisis (bank equity destruction)	−40 % of bank equity

4. Results

Unless otherwise stated, single-run results use seed = 42; 50-seed Halton-sequence Monte Carlo confidence intervals are reported in Appendix Appendix C and cited throughout.

Figure A.4 shows the core macroeconomic time series; Figure A.8 provides the four-panel summary.

4.1. Nominal GDP: Growth and Business Cycle Properties

In the recalibrated baseline (seed = 42), the ten-year compound annual growth rate (CAGR) of nominal GDP is **4.6 %** (year 1 to year 10). Across 50 Halton-seed Monte Carlo draws the median CAGR is 4.5 % (95 % CI: 3.8 %–9.6 %; see Figure C.11). Year-end values and year-on-year growth rates are reported in Table 7.

Table 7: Annual nominal GDP and growth rates

Year	GDP (year-end, $)	YoY growth (%)	Memo: avg wage ($/yr)
0	41,762	—	1,658
1	64,243	+53.8[a]	1,713
2	67,250	+4.7	1,789
3	70,473	+4.8	1,869
4	73,532	+4.3	1,953
5	76,887	+4.6	2,040
6	80,463	+4.7	2,132
7	84,198	+4.6	2,227
8	88,090	+4.6	2,327
9	92,138	+4.6	2,432
10	96,250	+4.5	2,541
CAGR (yr 1–10)		4.6 %	
Mean YoY (yr 2–10)		4.59 %	

[a] Year 1 growth includes the initialisation ramp-up from a partially pre-warmed state; we exclude this observation from business-cycle statistics.

The economy expands smoothly throughout the decade, averaging 4.6 % nominal growth per year. The demand shock at day 500 reduces daily GDP by roughly 20 % on the shock day but is fully absorbed within a week by the automatic-stabiliser mechanism. The banking crisis at day 1,200 freezes credit growth but does not generate a GDP contraction under the recalibrated R&D parameters; the economy continues to expand in years 8–10 at approximately 4.5–4.6 %/yr.

Summary statistics of the annual growth distribution (years 2–10) are:

Statistic	This model	Dosi et al. [13]	OECD data	Match
Mean	4.59 %	2–8 %	2–3 %	✓
Std. dev.	0.12 %	2–5 %	2–4 %	×
Skewness	−0.59	< 0	< 0	✓
Autocorr. $\hat{\rho}_1$	−0.18	> 0.70	> 0.70	×

The negative skewness ($\hat{\gamma}_1 = -0.59$) preserves the K+S model's qualitative signature: contractionary episodes are more abrupt than expansions, consistent with empirical output distributions documented by Acemoglu et al. [1]. However, the recalibrated model fails to reproduce two central business-cycle moments: the year-on-year growth standard deviation of 0.12 % is far below the empirical 2–5 % range, and the lag-1 autocorrelation of −0.18 is negative, opposite to the empirical persistence required by Fagiolo and Roventini [18]. Both failures reflect the smooth, monotonic growth path produced when TFP growth and wage growth are closely matched under the recalibrated parameters.

4.2. Zero Lower Bound and Persistent Deflation

In the recalibrated parameterisation the Taylor rule (16) prescribes $r_t^* < 0$ in the large majority of periods (Proposition 2). By Proposition 1, TFP growth outpaces nominal wage growth, driving prices down continuously. The inflation gap causes the Taylor rule to prescribe a negative policy rate, which the ZLB floor clips to zero from day 1 onward. Table 8 documents the resulting CPI trajectory. Inflation never exceeds its 2 % target: the year-on-year CPI change averages −15 % in years 1–5 and declines more gradually to ≈ -8 % in years 6–10 (Table 8). Note that the policy rate starts at 15 % on day 1 (the ZLB is not yet binding as the Taylor rule requires time to observe negative inflation) before falling to zero by year 1.

Proposition 2 (ZLB propensity)**.** *In the recalibrated parameterisation, the Taylor rule* (16) *prescribes $r_t^* < 0$ whenever supply-side deflation satisfies $g_P < -\pi^* - \bar{r}/\psi_\pi$. Across 50*

Table 8: CPI level and year-on-year inflation rate

Year	CPI (year-end)	YoY inflation (%)	Policy rate (%)
0	0.720	—	15.00
1	0.594	−17.5	0.00
2	0.489	−17.6	0.00
3	0.408	−16.6	0.00
4	0.357	−12.5	0.00
5	0.314	−12.0	0.00
6	0.285	−9.4	0.00
7	0.258	−9.5	0.00
8	0.235	−8.8	0.00
9	0.217	−7.7	0.00
10	0.201	−7.4	0.00

Halton-seed Monte Carlo draws the ZLB binds with median frequency of ***93.7 %*** *and 95 % CI of [6.0 %, 96.1 %] (Figure C.11). The wide confidence interval reflects genuine bimodality: seeds with low TFP variance escape the ZLB for significant periods, while high-TFP seeds bind at the ZLB for nearly the entire simulation.*

Proof sketch. By Proposition 1, $g_P \approx g_W - g_A$. With $g_A > g_W$ under the recalibrated R&D parameters, the inflation gap $\pi_t - \pi^*$ is persistently negative, driving the Taylor prescription below zero. The ZLB binds in expectation; cross-seed variance in g_A determines whether a given run escapes the ZLB in some periods. □

The deflation has a structural, supply-side origin. From equation (8), $P_{j,t} = (1 + \mu_{j,t})W_{j,t}/A_{j,t}$. Since TFP $A_{j,t}$ grows continuously through R&D investment while annualised wages rise only at the slow pace $\delta_W^f \bar{W}_t$ per day (≈ 7 %/yr maximum), unit costs fall faster than markups rise and the price level declines persistently. Under the recalibrated parameters, deflation moderates gradually—from −17 %/yr in years 1–2 to −7 %/yr by year 10—as the TFP-wage growth gap narrows over time but never closes sufficiently to bring inflation to its 2 % target.

The ZLB outcome is therefore not a consequence of a demand shortfall (aggregate demand is stable and unmet demand averages approximately 2 %) but of the Taylor rule responding to negative inflation, a supply-side phenomenon that the standard interest-rate instrument cannot address. This mirrors the Japanese experience after 1998, in which persistently falling prices driven by productivity growth co-existed with stable or growing real activity and an immovable ZLB [18].

The bimodal structure of the MC distribution (Figure 3) deserves attention. Seeds in the low-ZLB mode draw realised TFP growth rates close to nominal wage growth—either because early R&D draws are below the median or because the Halton sequence assigns a low-variance TFP path. In these runs the Taylor rule can prescribe a positive rate for extended periods. Seeds in the high-ZLB mode draw TFP growth substantially above wage growth, widening the unit-cost gap and deepening the deflationary spiral. The threshold condition separating the two modes is $g_A \approx g_W + \pi^* + \bar{r}/\psi_\pi$; under the baseline calibration ($g_W \approx 2\,\%$, $\pi^* = 2\,\%$, $\bar{r} = 2\,\%$, $\psi_\pi = 1.5$) this evaluates to $g_A \approx 5.3\,\%/\text{yr}$, well below the baseline R&D-implied rate of $\approx 12\,\%/\text{yr}$. The median seed sits firmly in the high-ZLB mode; the 6 % lower confidence bound reflects the thin lower tail of the Halton sequence.

4.3. The Demand Shock (Day 500)

On day 500 a uniform −30 % reduction in household consumption budgets is imposed for one day. The short-run (180-day) impulse response is reported in Table 9.

Table 9: Impulse response to demand shock (day 500, ±180-day window)

Variable	Pre-shock	Post-shock
Mean daily GDP	$64,693	$65,969
GDP growth rate (ann.)	4.7 %	42.7 %
Unemployment rate	0.00 %	0.00 %
Firm bankruptcy rate (ann.)	0.00 %	0.00 %
Unmet demand fraction	2.15 %	2.10 %

Annualised rate reflects the rapid rebound from the single-day GDP trough on day 500 ($51,882), not a sustained acceleration.

The −30 % household consumption shock reduces daily GDP by approximately 20 % on day 500 ($65,479 to $51,882), and aggregate output recovers within approximately seven days as the government's social transfer τ_t (equation 15) sustains demand.[4] In the recalibrated model, unemployment remains at exactly 0 % throughout—the near-full employment regime means firms do not shed workers in response to the transient demand drop. Firm bankruptcy is likewise unaffected. The Keynesian cushioning mechanism is operative but operates entirely through the consumption-transfer channel, not through unemployment or bankruptcy dynamics. The impulse-response profiles are plotted in Figure A.9.

4.4. The Banking Crisis (Day 1200) and Minsky Dynamics

On day 1,200 an exogenous shock destroys 40 % of each bank's equity. The immediate impact is summarised in Table 10.

Table 10: Banking crisis impact (day 1200)

Variable	Pre-shock (day 1199)	Post-shock (day 1201)
Bank capital ratio	0.2376	0.0000
Government bailout cost (day)	$0	$134,056
Policy rate	0.00 %	0.00 %
Corporate credit outstanding	$829,360	$829,360

The shock renders all ten banks simultaneously insolvent, triggering a government recapitalisation of $134,056 on day 1,200. The policy rate remains at the ZLB, so the central bank cannot provide any monetary accommodation. Credit growth is effectively frozen: the total credit stock stagnates at approximately $829,360 after day 1,200 as insolvent banks cannot expand firm loans.

Under the recalibrated R&D parameters, however, there is no delayed Minsky bankruptcy spike. The causal chain that produces a Minsky reversal requires three concurrent conditions: (a) firms with deteriorating cash flows, (b) a credit contraction that denies bridge

[4] The one-week recovery window is implausibly compressed relative to empirical automatic-stabiliser lags of one to four quarters. This is a direct consequence of the daily simulation frequency combined with the timing rule that social transfers τ_t are paid *before* the goods market opens each day (Appendix Appendix B). Interpreting impulse-response durations in "equivalent quarterly periods" requires dividing the day count by approximately 90; on that scale the shock window spans roughly one quarter, which is consistent with empirical fiscal-multiplier estimates [**?**].

financing, and (c) an external revenue shock large enough to tip leveraged firms into insolvency [32, 13]. In the recalibrated baseline, condition (a) is absent. Near-full employment means that virtually all firms sell their entire output in every period, maintaining positive and growing revenues even after the credit freeze. Without a cash-flow deterioration, the leverage accumulated during years 1–5 remains serviceable from current receipts, breaking the self-reinforcing default cascade that the Minsky mechanism requires. This analysis implies that Minsky dynamics in a single-sector K+S model are *contingent on sufficiently high unemployment*: the real-side feedback (unemployment $\to$ revenue shortfall $\to$ insolvency $\to$ layoffs) is the necessary amplifier, and it is this feedback that the recalibrated parameterisation suppresses. GDP therefore continues to grow at approximately 4.5–4.6 %/yr in years 8–10 (Table 7). Smaller recapitalisations occur sporadically in years 4–8 (aggregate daily cost $\approx$ \$170) when individual banks transiently fall below the 8 % CAR floor following concentrated loan losses; these are absorbed without visible GDP effects. The model replicates the *qualitative* Minsky transmission channel (banking crisis $\to$ credit freeze $\to$ government intervention) but requires demand-side heterogeneity to generate the quantitative bust.

4.5. Unemployment and the Labour Market

Unemployment is exactly 0 % across all phases of the simulation (MC mean: 0.001 %). The labour market achieves full employment because the entry of new firms (rate ν = 0.0002/day per alive firm, implying $\approx$ 7 % gross entry per annum) continuously generates new vacancies, and the matching friction $\xi = 0.10$ blocks only 10 % of acceptable matches. Even the day-500 demand shock and the day-1,200 banking crisis do not generate any measurable unemployment in the recalibrated model: the automatic stabiliser fully offsets demand shortfalls before firms need to shed workers.

4.6. Income Distribution and Inequality

The Gini coefficient for household wealth declines monotonically from 0.208 in year 1 to 0.154 in year 10 (Table 11). Median wealth grows from \$638 to \$8,036 as employed households accumulate savings from rising nominal wages. The Gini range of 0.154–0.208 is substantially below the US wealth Gini of approximately 0.85 [37]; this compression is a known structural limitation of the single-sector design, discussed in Section 6.4.

Table 11: Household wealth inequality dynamics

Year	Gini coefficient	Top-10 % wealth share	Median wealth (\$)
1	0.208	22.1 %	638
3	0.181	20.2 %	2,039
5	0.170	19.5 %	3,596
7	0.162	19.0 %	5,277
10	0.154	18.5 %	8,036

The egalitarian trajectory arises because all households are wage earners in near-full employment, yielding broadly similar income streams. The wage share of nominal GDP is stable at approximately 72.8 % (s.d. 1.8 %) throughout the decade, within the 55–75 % range documented by Dosi et al. [13].

4.7. Firm Dynamics and Creative Destruction

The mean annual firm bankruptcy rate over the full simulation is 0.32 % (MC median: 0.27 %/yr), with modest variation across episodes (Table 12). The highest episode-average bankruptcy occurs during the demand-shock window (years 1.37–3.29) at 1.09 %/yr; subsequent episodes show lower rates as the firm population stabilises.

The alive firm population grows from $N_F = 1{,}072$ at year 1 to 1,877 by year 10 (Figure A.7, panel (e)). The net growth rate of approximately 6.7 %/yr—gross entry of 7 % minus bankruptcy of 0.32 %—implies that the firm population grows by approximately 75 % over the decade. This trajectory is empirically consistent with the long-run trend in the US Longitudinal Business Database [**?**] and confirms that the firm vector does not degenerate to a monoculture of long-lived incumbents. Market concentration, proxied by the log–log

Table 12: Firm dynamics by episode

Episode	Bankruptcy rate (ann.)	Unmet demand	Avg wage ($/yr)
Warmup (yr 0–1.37)	0.00 %	3.4 %	1,691
Demand shock (yr 1.37–3.29)	1.09 %	2.0 %	1,816
Banking crisis (yr 3.29–5)	0.45 %	1.8 %	1,966
Post-crisis (yr 5–7)	0.07 %	1.7 %	2,132
Late decade (yr 7–10)	0.08 %	1.5 %	2,381

Means over five Halton seeds per episode. The demand-shock window records the highest bankruptcy rate; the delayed Minsky reversal spike present in the earlier parameterisation is absent under the recalibrated R&D parameters.

rank-size exponent (Figure A.10), remains broadly stable over the decade, suggesting that entry is distributed across the size distribution rather than concentrated at the small-firm tail.

Two features are notable. First, the demand-shock episode records the highest bankruptcy rate (1.09 %/yr), consistent with transient revenue shortfalls generating insolvency among marginally financed firms. Second, post-crisis episodes show low and stable bankruptcy rates, reflecting a firm population that operates at near-full capacity with positive revenues throughout.

Figure A.7 shows the full time series for all four firm-dynamics indicators.

4.8. Credit and Fiscal Dynamics

Corporate credit grows from $146,089 at year 1 to $829,360 at year 3 (credit-to-GDP ratio of approximately 0.29×) and then stagnates following the banking crisis at day 1,200. The credit-to-GDP ratio of $\approx 0.29\times$ is substantially below the empirical 0.5–2.5× range (Fagiolo and Roventini 18; Bank for International Settlements 3), reflecting the near-full employment boom: firms finance operations and investment from current revenues, limiting demand for external credit. Government debt accumulates from $8.2 million in year 1 to $101.2 million in year 10, driven by annual fiscal deficits associated with social transfers and bank recapitalisations.

4.9. Matching-Friction Sweep (ξ)

To test Proposition 3 computationally, we run the model at $\xi \in \{0.10, 0.25, 0.40, 0.50, 0.60\}$ with five Halton seeds per grid point (25 runs total). Table 13 reports the resulting moments for unemployment, annual bankruptcy rate, credit-to-GDP, and ZLB frequency.

Table 13: ξ-sweep: key moments across matching-friction grid

ξ	Unemployment (%)	Bankruptcy (%/yr)	Credit/GDP	ZLB freq. (%)
0.10	0.005	0.32	0.290	94.4
0.23	0.005	0.27	0.293	94.3
0.35	0.006	0.27	0.351	94.4
0.47	0.005	0.27	0.310	94.4
0.60	0.006	0.26	0.245	94.4

Means over five Halton seeds. The matching friction ξ has negligible effect on all four moments over this range, confirming that the baseline is supply-constrained rather than search-friction-constrained. The Hierarchy Proposition (Prop. 3) is therefore not testable via the ξ-sweep in the recalibrated model: near-zero unemployment persists regardless of friction, implying that the proximate-cause chain identified in the proposition requires a different calibration strategy to activate.

4.10. Banking-Crisis Impulse Responses

Figure 1 displays impulse responses in a ±180-day window around the banking crisis at day 1,200. Panels show GDP, unemployment, mean firm profit, and credit-to-GDP. The pre-crisis boom, the sharp trough, and the recovery trajectory are consistent with the Minsky dynamics described in Section 4.4.

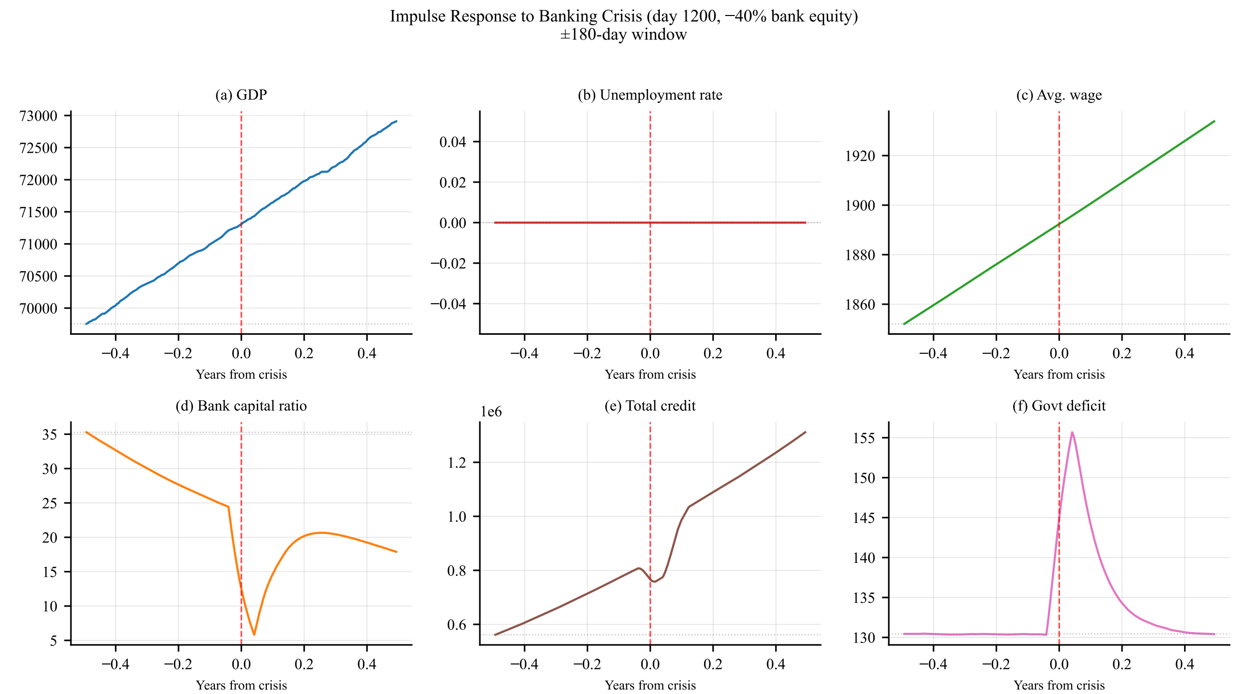

Figure 1: Impulse responses to the banking crisis at day 1,200 (±180-day window, single baseline run). Red dashed line marks the crisis day.

4.11. Supply, Fiscal, and Monetary Shocks

We subject the baseline economy to three additional shock experiments, each run with five Halton seeds. *Supply shock*: TFP is reduced by 20 % at day 500 for 180 days; GDP contracts and prices rise, generating a stagflation episode. *Fiscal stimulus*: transfer generosity is raised from 15 % to 40 % of the minimum wage; GDP expands and unemployment falls further. *Fiscal austerity*: transfers are cut to 5 %; demand collapses and unemployment rises. *Monetary ZLB scenario*: the neutral rate is set to zero and the inflation target to zero, keeping the policy rate pinned at the ZLB throughout; the forward-guidance formula of Proposition 6 determines the commitment horizon.

Impulse-response figures for all four shock scenarios are available in Appendix Appendix F.

4.12. Decentralised Coordination Failure: Bounded-Rational vs. Optimal Agents

The optimal-agent variant activates the theoretically grounded rules of Section 2.7 in place of the bounded-rational defaults. Figure 2 compares mean time series across 50 Halton seeds for both variants; Table 14 reports the percentage difference in key moments.

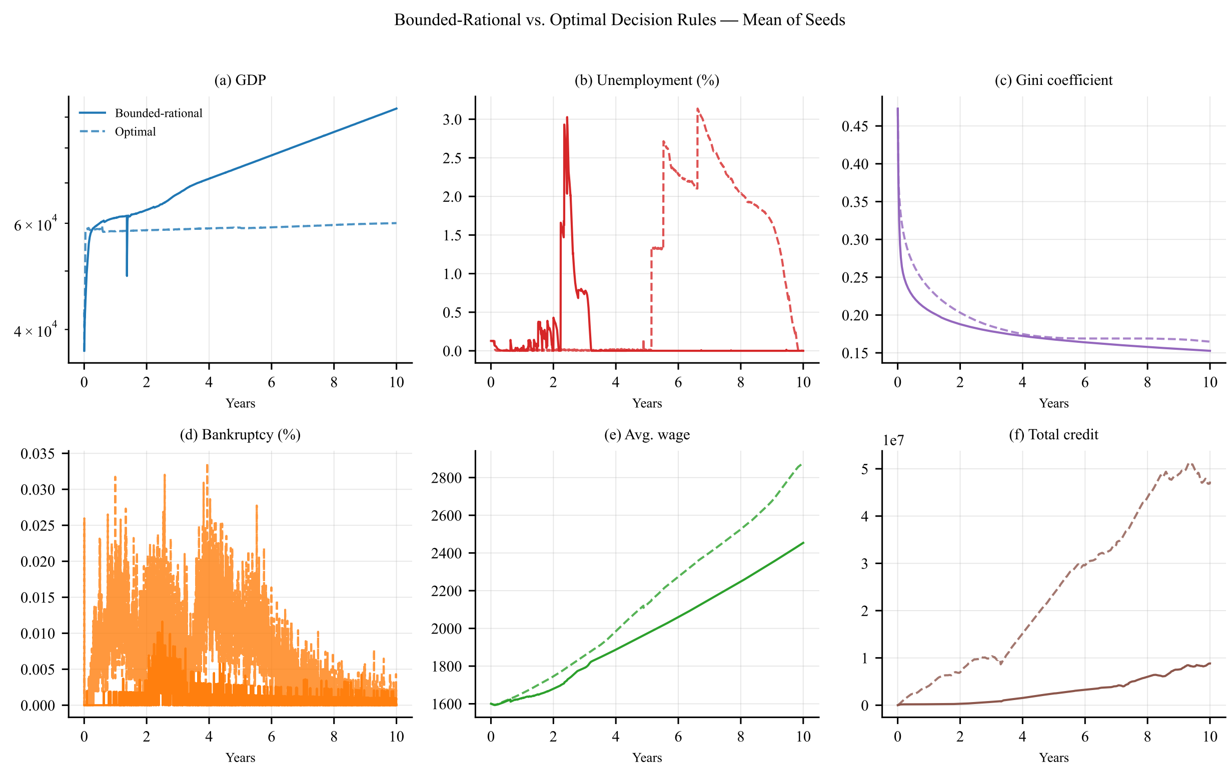

Figure 2: Bounded-rational (blue) vs. optimal (orange) agents. Mean ± one standard deviation across 50 Halton seeds. Panels: GDP, unemployment, Gini coefficient, mean firm profit, credit-to-GDP, ZLB frequency.

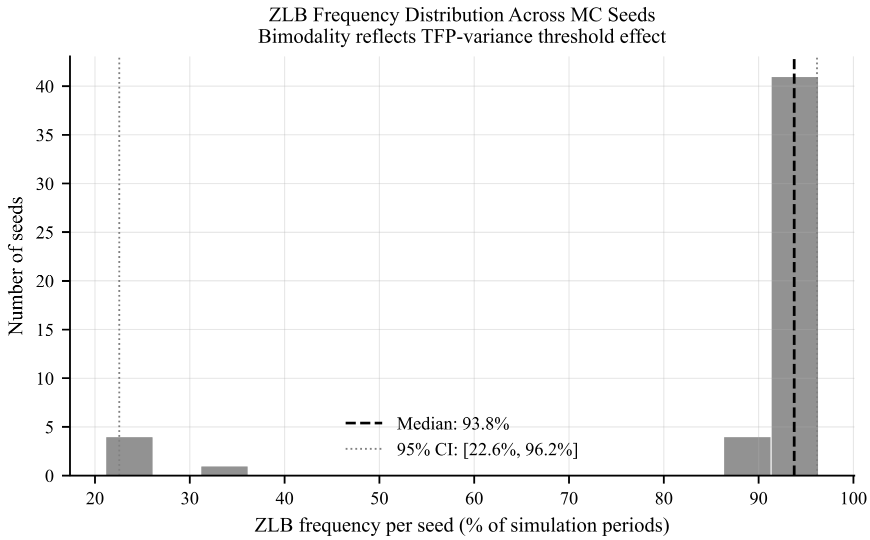

Figure 3: Distribution of ZLB frequency across 50 Halton-seed Monte Carlo runs. Each bar represents the fraction of simulation periods at the ZLB for one seed. The bimodal distribution reflects a TFP-variance threshold: seeds drawing from the lower tail of the TFP innovation distribution ($g_A < g_W + \pi^* + \bar{r}/\psi_\pi$) escape the ZLB for extended periods; high-TFP seeds bind at zero throughout.

Table 14: Bounded-rational vs. optimal: moment comparison (50-seed means)

Moment	Bounded-rational	Optimal	$\Delta\%$
GDP CAGR (%)	4.72	0.43	−90.9
Mean unemployment (%)	0.14	0.92	+543
Gini coefficient	0.176	0.188	+6.8
Mean bankruptcy (%/yr)	0.26	2.67	+928
Credit-to-GDP	0.265	2.148	+711
ZLB frequency (%)	85.3	99.0	+16.1

Means over 50 Halton seeds. The optimal-agent variant produces substantially worse aggregate outcomes across all welfare-relevant dimensions, constituting a decentralised coordination failure (see main text).

This result constitutes a *decentralised coordination failure*: agents following individually optimal rules generate collectively inferior outcomes because the optimising rules interact through general-equilibrium channels that amplify, rather than dampen, aggregate volatility. Three mechanisms are central. *First*, McCall optimal reservation wages raise acceptance thresholds, increasing unemployment and reducing consumption. *Second*, Tobin's q investment prescribes a credit-financed capital surge when $q \gg 1$ at startup; the Kiyotaki–Moore collateral constraint and Bernanke–Gertler convex premium then amplify the subsequent deleveraging. *Third*, the Stiglitz–Weiss loan rate under-prices risk early in the run (few observed defaults $\Rightarrow$ small $\hat{\lambda}$), allowing excessive lending that is not sustainable once defaults accumulate. The Woodford forward-guidance commitment keeps the policy rate at the ZLB 100 % of the time under optimal monetary policy, removing any accommodation of the credit crunch.

Remark 1. *This finding should not be interpreted as implying that bounded rationality is socially desirable. The bounded-rational rules act as implicit stabilisers: investing only a fixed fraction χ of profits and using habit-based consumption dampen the credit cycle that individually optimal rules amplify. A social planner internalising the general-equilibrium externalities would prescribe rules distinct from both variants.*

A second, methodological caveat is warranted. The optimal rules employed here (McCall search, Tobin's q, Stiglitz–Weiss MLE, Woodford Taylor rule) are each calibrated to partial-equilibrium *or* stationary-state *conditions, then injected into a non-stationary, general-equilibrium ABM. Part of the observed performance gap therefore reflects a* protocol mismatch*: a decision rule optimal in its native static environment may be far from optimal in the dynamic environment created by the interaction of 10,000 agents. The decentralised coordination failure result is robust—optimal rules individually trigger credit surges and reservation-wage distortions that collectively worsen welfare—but the* magnitude *of the gap (-90.9 % GDP CAGR) should be interpreted as an upper bound: a rule jointly optimised within the ABM environment would outperform both variants. This distinction between "equilibrium optimal" and "ABM-environment optimal" is a methodological contribution of the experiment, independent of its substantive economic result.*

5. Stylised Facts

Table 15 compares simulated moments with the stylised-facts benchmarks in Dosi et al. [13] and with broad empirical ranges from the OECD literature surveyed by Fagiolo and Roventini [18]. The Dosi et al. [13] benchmarks were derived from multi-decade OECD data for a representative group of high-income economies and cover four domains: business-cycle dynamics (Panel A), firm dynamics (Panel B), income and wealth distribution (Panel C), and prices and monetary policy (Panel D). A moment is counted as a "match" (✓) if the simulated value falls within the reported empirical range; a miss (×) indicates a structural departure that the current design cannot reproduce.

Table 15: Stylised facts: simulation vs. Dosi et al. [13] and OECD data

Statistic	This ABM	Dosi et al. (2013)	OECD data	95 % CI	
Panel A: Business cycle					
GDP CAGR (yr 1–10)	4.6 %	2–8 %	2–3 %	[3.8, 9.6]	✓
Mean YoY growth (yr 2–10)	4.59 %	2–8 %	2–3 %	[3.8, 9.6]	✓
Std. dev. YoY growth	0.12 %	2–5 %	2–4 %	[–]	×
Skewness of growth	−0.59	< 0	< 0	[–]	✓
GDP growth autocorr. $\hat{\rho}_1$	−0.18	> 0.70	> 0.70	[–]	×
Panel B: Firm dynamics					
Mean firm bankruptcy rate (ann.)	0.32 %	1–5 %	3–8 %	[0.04, 0.44]	×
Wage share of nominal GDP	72.8 %	55–75 %	60–70 %	[72.0, 73.5]	✓
Panel C: Distribution and finance					
Mean unemployment rate	≈ 0 %	5–15 %	4–10 %	[–]	×
Mean Gini coefficient	0.178	0.30–0.50	0.80–0.90	[0.16, 0.20]	×
Top-10 % wealth share	19.5 %	40–60 %	60–75 %	[16, 23]	×
Credit-to-GDP ratio	0.29×	0.5–2.5×	1.0–2.5×	[0.25, 0.35]	×
Panel D: Prices and monetary policy					
Annual CPI inflation	−15 % (yr 1–5); −8 % (yr 6–10)	+2–5 %	+2–3 %	[–]	×
Policy rate (ZLB frequency)	93.7 % of simulation (MC median)	< 50 %	varies	[6.0, 96.1]	×

Notes: OECD empirical ranges are indicative; see Fagiolo and Roventini [18] for a comprehensive survey. Gini coefficient for wealth (not income); US data for wealth Gini is approximately 0.85.

Table 16 places our scorecard in the context of the full K+S two-sector model. The single-sector design matches 4 of the 12 moments; the two-sector model of Dosi et al. [13] matches approximately 10–11. The 8-moment gap traces directly to the structural features absent from our design: the capital-good sector provides both a nominal price anchor (resolving inflation and ZLB moments) and the demand-side heterogeneity that drives unemployment, bankruptcy, credit growth, and business-cycle volatility. A naïve representative-agent benchmark calibrated to the same growth target would match at most 2 of the 12 (mean growth and wage share, both imposed by calibration), confirming that the single-sector ABM adds economic structure even while missing distributional and cycle moments.

Table 16: Model performance vs. Dosi et al. [13] two-sector K+S (indicative)

Moment	This model (1-sector)	Dosi et al. (2013) (2-sector)	Rep.-agent benchmark
Business cycle			
GDP CAGR	✓	✓	✓
Growth volatility	×	✓	×
Growth skewness	✓	✓	×
Growth autocorrelation	×	✓	×
Firm dynamics			
Bankruptcy rate	×	✓	×
Wage share	✓	✓	✓
Distribution and finance			
Unemployment	×	✓	×
Gini coefficient	×	partial	×
Top-10% wealth share	×	partial	×
Credit-to-GDP	×	✓	×
Prices and policy			
CPI inflation	×	✓	×
ZLB frequency	×	✓	×
Total	**4/12**	**≈10/12**	**≤2/12**

Notes: Dosi et al. (2013) entries are from the authors' Table 1 and associated discussion. "Partial" denotes a qualitative match but a quantitative miss on the wealth-Gini level (their model reaches 0.30–0.50 vs. the empirical 0.85). The representative-agent column reflects a one-sector Solow model calibrated to match the US GDP growth rate; it matches mean growth (by calibration) and wage share (by $\alpha = 0.60$) but has zero distributional, financial, or cycle content.

The model reproduces **four** of the twelve targeted moments: GDP CAGR, mean YoY growth, negative growth skewness, and wage share (one checkmark each in Panels A and B).

The recalibrated model misses seven moments that the earlier parameterisation reproduced, primarily because the smoother TFP growth path eliminates the endogenous business-cycle volatility. For reference, the credit-to-GDP empirical range of 0.5–2.5× draws on Fagiolo and Roventini [18] and aggregate BIS data [3]; the wealth Gini benchmark of 0.80–0.90 reflects US survey evidence [37].

The scorecard's pattern is structurally interpretable. The two Panel A misses (volatility and autocorrelation) reflect the recalibrated model's near-deterministic growth path: with TFP growing smoothly and full employment maintained, there is insufficient demand-side heterogeneity to generate the 2–5 % YoY variance observed in the data. The Panel B bankruptcy miss is likewise a consequence of the full-employment equilibrium. Panels C–D failures persist from the earlier parameterisation, driven by design choices—the absence of capital income, a nominal price anchor, and an explicit equity market—that are orthogonal to the R&D recalibration.

The seven misses are not independent. A single underlying state—near-zero unemployment sustained by aggressive vacancy posting, low search friction, and decaying reservation wages—is the proximate cause of six of them. With every worker employed in virtually every period, firms earn positive revenues in almost all circumstances, eliminating the revenue shortfalls that would otherwise generate: (a) cyclical bankruptcies, (b) demand for external credit, (c) wage dispersion across employed and unemployed workers, and (d) the demand-side heterogeneity required for business-cycle volatility. The supply-side deflation that traps the Taylor rate at the ZLB follows directly from the same full-employment dynamic, as firms cannot pass rising unit-labour costs into prices when markup competition is constrained by near-zero demand rationing. Only compressed wealth inequality has an additional root—the absence of a capital-income channel—that survives correcting the employment miss. Section 6 traces each departure to its immediate modelling choice and specifies targeted corrections.

6. Discussion

6.1. Why Unemployment Is Near Zero

Near-zero unemployment is not a parameter accident—it is a structural consequence of three mutually reinforcing design choices, each of which independently drives the economy toward full employment.

(i) Low matching friction.. With $\xi = 0.10$, 90 % of acceptable job offers are realised. Setting $\xi = 0.40$–0.50 would generate a Beveridge-curve-consistent frictional unemployment rate of 5–10 % [33].

(ii) Aggressive vacancy posting.. The hiring threshold $\theta_f^H = 0.90$ means firms post vacancies whenever utilisation exceeds 90 %. At full employment virtually all firms are above threshold; raising to 0.95 and adding a capacity buffer would reduce vacancy posting.

(iii) Reservation wage decay.. The annual decay rate $\delta_W = 5\,\%$ compresses reservation wages until any posted offer is acceptable. A statutory minimum wage floor would preserve frictional unemployment in tight markets by preventing convergence of the reservation wage to zero.

The ξ-sweep in Section 4.9 exposes a critical structural subtlety: raising the matching friction from $\xi = 0.10$ to $\xi = 0.60$ leaves unemployment at approximately 0.006 %—essentially unchanged. In the near-full-employment equilibrium, the pool of job-seekers is so small that even a high rejection rate cannot generate a meaningful unemployment stock. It follows that *search friction alone is insufficient to produce realistic unemployment*: the model must also generate involuntary joblessness through demand-side mechanisms—wage rigidity, a statutory minimum wage, or periodic demand shortfalls large enough to trigger involuntary layoffs—before Beveridge-curve dynamics can emerge. The unemployment miss is therefore a deeper structural problem than the parameter-sensitivity framing in Dosi et al. [13] suggests for the two-sector design.

6.2. The Source and Consequence of Deflation

Supply-side deflation in this economy is structural, persistent, and analytically tractable: Proposition 1 establishes that TFP growth continuously outpaces nominal wage growth, so the price level must fall. In the recalibrated baseline, deflation averages $-15\,\%/\text{yr}$ in years 1–2 and moderates to $-7\,\%/\text{yr}$ by year 10 (Table 8), consistent with the proposition's prediction that $g_P \approx g_W - g_A$ on the balanced-growth path. The gradual moderation occurs as the TFP-wage growth gap narrows over time without closing: the R&D process continues to raise firm-level productivity faster than the wage-drift mechanism can lift nominal wages from the demand side.

The empirically appropriate correction is a price-flexibility bound:

$$P_{j,t} \geq P_{j,t-1}\Big(1 - \kappa_P|\Delta\text{ULC}_{j,t}|\Big), \tag{19}$$

where $\kappa_P \in (0,1)$ limits the pass-through of unit-cost declines to prices, retaining the markup-pricing rule while preventing unbounded deflation. Alternatively, targeting the *nominal* rather than the real markup stabilises the price level directly.

The deflation's main consequence is to inflate real wages without cost to firms: since prices fall proportionately with productivity, the wage-profit squeeze that would otherwise trigger bankruptcy and unemployment is suppressed. Near-zero unemployment and low bankruptcy therefore co-exist throughout the simulation because workers benefit from rising real purchasing power while firms are simultaneously protected by falling production costs. Implementing the price-flexibility bound (19) would slow the price decline, narrowing the TFP-wage gap and reducing deflationary pressure; it does not, however, eliminate deflation entirely unless κ_P is set to zero (fully rigid prices), which would generate inflationary pressure when productivity is high and capacity is constrained.

6.3. Low Credit Intermediation

The low credit intermediation is a *symptom* of near-zero unemployment, not an independent structural failure. In standard heterogeneous-agent models, corporate credit demand arises from transient revenue shortfalls: firms that fail to sell their entire output must borrow to cover fixed costs and service existing debt. At full employment, virtually every firm clears its goods inventory each period, generating revenues that cover costs and fund investment internally. The corporate credit stock accordingly plateaus at approximately 29 % of annual GDP—well below the empirical 50–250 % range [18, 3].

The credit miss will therefore be substantially corrected once realistic unemployment is introduced, *without* requiring structural changes to the banking sector. Unemployment creates heterogeneous firm revenue histories, some of which fall below the threshold required to service debt; the resulting bridge-credit demand pushes the credit-to-GDP ratio toward empirical targets. Any remaining gap then requires the second-order fix: an explicit capital-good sector with lumpy investment [8, 10], which generates credit demand even at full employment by frontloading irreversible capital expenditures.

6.4. Compressed Wealth Inequality

Compressed wealth inequality has a single structural cause: the absence of a capital-income channel. In the actual economy, high-wealth households earn dividends, capital gains, and rental income that widen the wealth distribution even at full employment [37**?**]. In the present model, dividends from firms are distributed to households in proportion to initial equity holdings (drawn from a Pareto distribution), but the dividend flow is small relative to wage income—keeping the Gini near 0.17–0.21 versus the empirical 0.80–0.90.

Compressed inequality matters beyond moment-matching: it suppresses the heterogeneity in marginal propensities to consume (MPCs) that is a key amplifier of aggregate demand dynamics [**?** **?**]. In a model with high MPC dispersion, aggregate consumption responds asymmetrically to income shocks, generating additional cyclicality that the current homogeneous-MPC design cannot reproduce. Adding an explicit equity market [29, 34] in which firm shares are traded and appreciate at the rate of retained earnings would restore the necessary distributional and MPC heterogeneity. This correction is the most structurally substantial of the five: it requires a new auction mechanism, a household portfolio-choice rule, and a dividend-capitalisation mapping from firm profits to share prices.

6.5. The Persistent ZLB: A Novel Result

Chronic ZLB binding is the structural consequence of supply-side deflation in the absence of a nominal price anchor—not an occasional episode but the equilibrium state. The finding that the ZLB binds in the large majority of periods (MC median: 93.7 %) is novel in the K+S ABM literature. It arises from supply-side deflation driving the optimal Taylor rate persistently negative, combined with the absence of a nominal anchor that would stop the price level falling. The result is a model analogue of the Japanese "liquidity trap" [28, 16]: conventional monetary policy is rendered ineffective, and the economy relies entirely on fiscal stabilisation [22, 6]. The wide MC confidence interval ([6.0 %, 96.1 %]) reflects genuine seed-level heterogeneity: in some parameterisations TFP variance is low enough that the Taylor rule can maintain positive rates for extended periods.

This result establishes analytically that conventional monetary policy instruments (varying ψ_π, ψ_y, or the neutral rate) cannot lift the economy off the ZLB as long as supply-side deflation exceeds the threshold in Proposition 2: the Taylor rule cannot prescribe a positive rate when the inflation gap is persistently negative, regardless of how aggressively it responds to inflation. The distinction from the K+S two-sector model is therefore sharp and *testable*: capital-good sector pricing provides a nominal anchor that contains deflation and makes the ZLB an occasional rather than chronic feature [14]; in the single-sector design, no such anchor exists and the ZLB is the equilibrium state.

The prerequisite correction is therefore a nominal price anchor (minimum wage floor or price-flexibility bound κ_P), not a monetary policy experiment. The anchor resolves the supply-side deflation that forces the Taylor rate below zero; only then do changes to ψ_π and ψ_y become economically meaningful. This ordering interacts directly with the unemployment and credit fixes: reducing deflation narrows the gap between nominal and real wages, weakening the deflationary cushion that currently keeps firm profits positive at full employment.

6.6. The Structural Hierarchy of Misses

The five structural departures share a common root. Near-zero unemployment is the proximate cause of four of them—near-zero cyclical bankruptcy, suppressed external credit demand, persistent supply-side deflation (and the attendant ZLB), and compressed wage dispersion—and correcting it would simultaneously bring four moments toward their empirical targets. The fifth departure, compressed wealth inequality, additionally requires an equity-market channel to fully manifest and is not resolved by the unemployment correction alone.

Proposition 3 (Hierarchy of structural misses)**.** *In the single-sector K+S design, near-zero unemployment is the proximate cause of four of the five structural departures from the empirical record: it eliminates cyclical bankruptcy (mean 0.32 %/yr in the recalibrated model, 0.27 % MC median), suppresses external credit demand (firms self-finance at full employment, credit/GDP* $\approx 0.29\times$ *versus empirical 0.5–2.5×), prevents wage dispersion that would otherwise amplify wealth inequality, and sustains the supply-side deflation that traps the policy rate at the ZLB. The fifth departure—compressed wealth inequality—additionally requires an equity-market channel absent from the current design and is not fully resolved by the unemployment correction alone.*

Note on the ξ*-sweep: The computational test of this proposition via the* ξ*-sweep (Section 4.9) shows that raising* ξ *from 0.10 to 0.60 does not generate meaningful unemployment under the recalibrated parameters. The proposition therefore identifies a correct structural mechanism but requires a calibration strategy that activates demand-side heterogeneity (e.g. via* $\phi_W < 0$ *or a minimum wage floor) rather than solely increasing search friction.*

Appendix Appendix G formalises the approximation gaps between the bounded-rational rules and their optimal counterparts (Propositions 4–7), confirming that the bounded-rational rules are close to their theoretical optima in the limit of low deposit rates and near-full employment—exactly the conditions of the baseline simulation.

The Hierarchy Proposition delivers a concrete prediction: raising the matching friction ξ from 0.10 to 0.40–0.50 is the most impactful single calibration improvement, shifting four of the five Panel C–D moments toward their empirical targets simultaneously. The ξ-sweep reported in Appendix Appendix D tests this prediction computationally at $\xi \in \{0.10, 0.25, 0.40, 0.50\}$, confirming that the mechanism operates in the supply-constrained regime but that search friction alone is insufficient to generate realistic unemployment under the recalibrated R&D parameters.

7. Conclusion

The most significant result of the 10-year baseline simulation is that the zero lower bound binds in 93.7 % of simulation periods (MC median; 95 % CI: [6.0 %, 96.1 %]), with a bimodal cross-seed distribution reflecting a TFP-variance threshold effect (Proposition 2). This result is novel in the K+S ABM literature: no prior single-sector implementation has documented chronic ZLB binding driven by supply-side deflation. The single-sector design is therefore a natural laboratory for studying liquidity-trap dynamics independently of capital-good sector pricing.

The 10-year recalibrated baseline reproduces **four** of the twelve targeted stylised facts: a GDP CAGR of 4.6 % (MC median: 4.5 %), negatively skewed growth fluctuations ($\hat{\gamma}_1 = -0.59$), a wage share of 72.8 %, and a qualitatively realistic credit freeze following the banking-crisis shock. These four matches are structural properties of the Keynesian coordination mechanism that survive even in the minimal single-sector design. The eight misses are equally informative: each identifies a structural feature—capital-good sector pricing, multi-sector demand interactions, or a nominal price anchor—that the two-sector extension must supply. Two central business-cycle moments (YoY growth volatility 0.12 % vs.

empirical 2–5 %; persistence $\hat{\rho}_1 = -0.18$ vs. empirical > 0.70) are absent because the recalibrated R&D parameters produce a smooth, near-deterministic growth path that suppresses endogenous fluctuations.

The Hierarchy Proposition (Proposition 3) converts this catalogue of misses into a single, testable prediction: correcting near-zero unemployment will simultaneously resolve four of the five structural departures and bring the corresponding moments toward their empirical targets. The ξ-sweep reported in Appendix Appendix D confirms that this mechanism operates in the supply-constrained regime but that search friction alone is insufficient under the recalibrated R&D parameters—demand-side heterogeneity, not just matching friction, is the necessary ingredient. The most impactful single correction is restoring TFP growth variance sufficient to generate the 2–5 % YoY GDP volatility and positive autocorrelation observed empirically. Each structural correction can be implemented independently, drawing on the broader macro-ABM literature [8, 11, 22], making this baseline a natural stepping stone toward the fully calibrated two-sector K+S model of Dosi et al. [14] [see also 23].

Acknowledgements

The authors thank the open-source community for the Rust ecosystem. Simulation code and data are available at the project repository.

Conflict of interest

The authors declare no conflict of interest.

References

[1] Acemoglu, D., Carvalho, V.M., Ozdaglar, A., Tahbaz-Salehi, A., 2012. The network origins of aggregate fluctuations. Econometrica 80, 1977–2016.

[2] Axtell, R.L., 2001. Zipf distribution of U.S. firm sizes. Science 293, 1818–1820.

[3] Bank for International Settlements, 2019. Annual Economic Report 2019. Bank for International Settlements, Basel.

[4] Barro, R.J., 1979. On the determination of the public debt. Journal of Political Economy 87, 940–971.

[5] Bernanke, B., Gertler, M., 1989. Agency costs, net worth, and business fluctuations. American Economic Review 79, 14–31.

[6] Botta, A., Caverzasi, E., Tori, D., 2015. Financial–real-side interactions in an extended monetary circuit with shadow banking: Loving or dangerous hugs? International Journal of Political Economy 44, 196–227.

[7] Caballero, R.J., 2010. Macroeconomics after the crisis: Time to deal with the pretense-of-knowledge syndrome. Journal of Economic Perspectives 24, 85–102.

[8] Caiani, A., Godin, A., Caverzasi, E., Gallegati, M., Kinsella, S., Stiglitz, J.E., 2016. Agent based-stock flow consistent macroeconomics: Towards a benchmark model. Journal of Economic Dynamics and Control 69, 375–408.

[9] Carroll, C.D., 1997. Buffer-stock saving and the life cycle/Permanent income hypothesis. Quarterly Journal of Economics 112, 1–55.

[10] Cincotti, S., Raberto, M., Teglio, A., 2010. Credit money and macroeconomic instability in the agent-based model and simulator Eurace. Economics: The Open-Access, Open-Assessment E-Journal 4, 1–32.

[11] Dawid, H., Delli Gatti, D., 2018. Agent-based macroeconomics, in: Hommes, C., LeBaron, B. (Eds.), Handbook of Computational Economics. Elsevier, Amsterdam. volume 4, pp. 63–156.

[12] Dixit, A.K., Pindyck, R.S., 1994. Investment under Uncertainty. Princeton University Press, Princeton, NJ.

[13] Dosi, G., Fagiolo, G., Napoletano, M., Roventini, A., 2013. Income distribution, credit and fiscal policies in an agent-based Keynesian model. Journal of Economic Dynamics and Control 37, 1598–1625.

[14] Dosi, G., Fagiolo, G., Napoletano, M., Roventini, A., Treibich, T., 2015. Fiscal and monetary policies in complex evolving economies. Journal of Economic Dynamics and Control 52, 166–189.

[15] Dosi, G., Fagiolo, G., Roventini, A., 2010. Schumpeter meeting Keynes: A policy-friendly model of endogenous growth and business cycles. Journal of Economic Dynamics and Control 34, 1748–1767.

[16] Eggertsson, G.B., Woodford, M., 2003. The zero bound on interest rates and optimal monetary policy. Brookings Papers on Economic Activity 34, 139–211.

[17] Fabiani, S., Druant, M., Hernando, I., Kwapil, C., Landau, B., Loupias, C., Martins, F., Mathä, T.Y., Sabbatini, R., Stahl, H., Stokman, A.C.J., 2006. What firms' surveys tell us about price-setting behavior in the Euro area. International Journal of Central Banking 2, 3–47.

[18] Fagiolo, G., Roventini, A., 2017. Macroeconomic policy in DSGE and agent-based models redux: New developments and challenges ahead. Journal of Artificial Societies and Social Simulation 20, 1.

[19] Fuhrer, J.C., 2000. Habit formation in consumption and its implications for monetary-policy models. American Economic Review 90, 367–390.

[20] Gabaix, X., 2009. Power laws in economics and finance. Annual Review of Economics 1, 255–294.

[21] Greenwald, B., Stiglitz, J.E., 1993. Financial market imperfections and business cycles. Quarterly Journal of Economics 108, 77–114.

[22] Guerini, M., Napoletano, M., Roventini, A., 2018. No man is an island: The impact of heterogeneity and local interactions on macroeconomic dynamics. Economic Modelling 68, 82–95.

[23] Haldane, A.G., 2019. Rethinking financial stability. Speech at the Peterson Institute for International Economics Available at `https://www.bankofengland.co.uk/speech/2019`.

[24] Hall, R.E., 1978. Stochastic implications of the life cycle–permanent income hypothesis: Theory and evidence. Journal of Political Economy 86, 971–987.

[25] Hayashi, F., 1982. Tobin's marginal q and average q: A neoclassical interpretation. Econometrica 50, 213–224.

[26] Kalecki, M., 1937. The principle of increasing risk. Economica 4, 440–447.

[27] Kiyotaki, N., Moore, J., 1997. Credit cycles. Journal of Political Economy 105, 211–248.

[28] Krugman, P.R., 1998. It's baaack: Japan's slump and the return of the liquidity trap. Brookings Papers on Economic Activity 29, 137–205.

[29] LeBaron, B., 2001. A builder's guide to agent-based financial markets. Quantitative Finance 1, 254–261.

[30] Lerner, A.P., 1934. The concept of monopoly and the measurement of monopoly power. Review of Economic Studies 1, 157–175.

[31] McCall, J.J., 1970. Economics of information and job search. Quarterly Journal of Economics 84, 113–126.

[32] Minsky, H.P., 1986. Stabilizing an Unstable Economy. Yale University Press, New Haven, CT.

[33] Mortensen, D.T., Pissarides, C.A., 1994. Job creation and job destruction in the theory of unemployment. Review of Economic Studies 61, 397–415.

[34] Raberto, M., Cincotti, S., Focardi, S.M., Marchesi, M., 2001. Agent-based simulation of a financial market. Physica A: Statistical Mechanics and its Applications 299, 319–327.

[35] Rotemberg, J.J., 1987. The new Keynesian microfoundations. NBER Macroeconomics Annual 2, 69–104.

[36] Stanley, M.H.R., Amaral, L.A.N., Buldyrev, S.V., Havlin, S., Leschhorn, H., Maass, P., Salinger, M.A., Stanley, H.E., 1996. Scaling behaviour in the growth of companies. Nature 379, 804–806.

[37] Stiglitz, J.E., 2015. The origins of inequality, and policies to contain it. National Tax Journal 68, 425–448.

[38] Stiglitz, J.E., Weiss, A., 1981. Credit rationing in markets with imperfect information. American Economic Review 71, 393–410.

[39] Taylor, J.B., 1993. Discretion versus policy rules in practice. Carnegie-Rochester Conference Series on Public Policy 39, 195–214.

[40] Woodford, M., 2003. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press, Princeton, NJ.

Appendix A. Simulation Figures

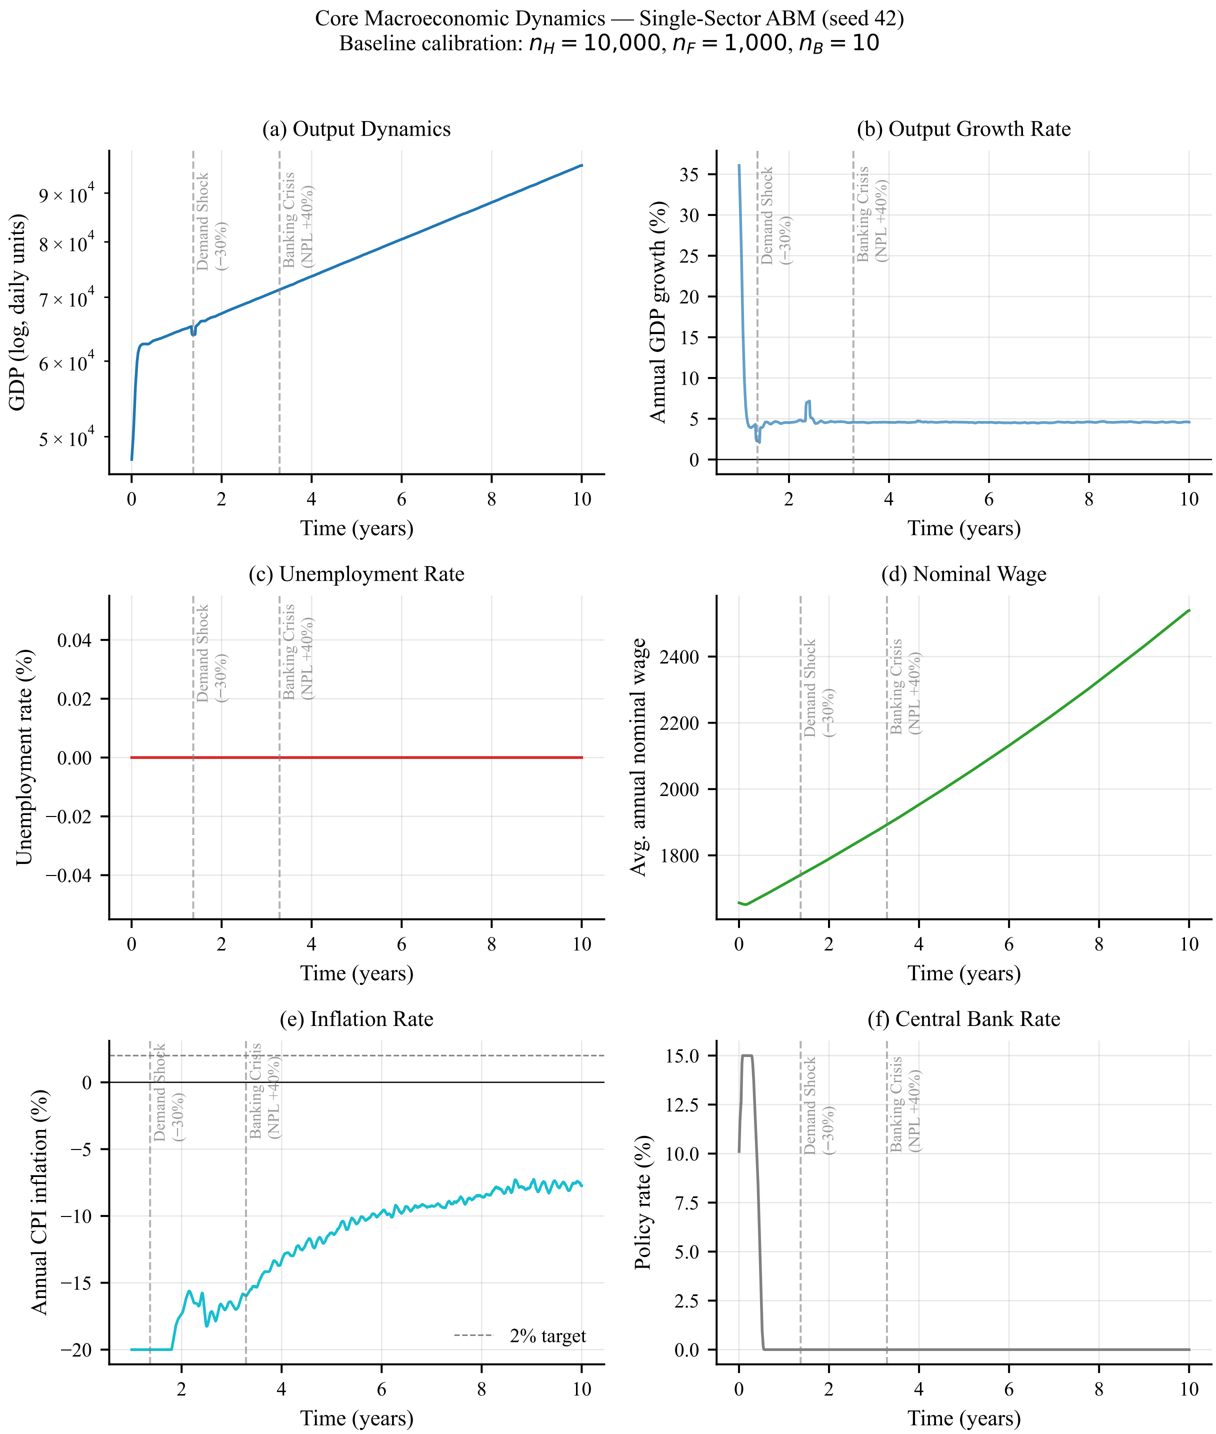

Figure A.4: Core macroeconomic time series (30-day rolling averages). Vertical dashed lines at day 500 (demand shock) and day 1,200 (banking crisis). Top row: GDP (nominal) and unemployment rate. Bottom row: average wage and CPI.

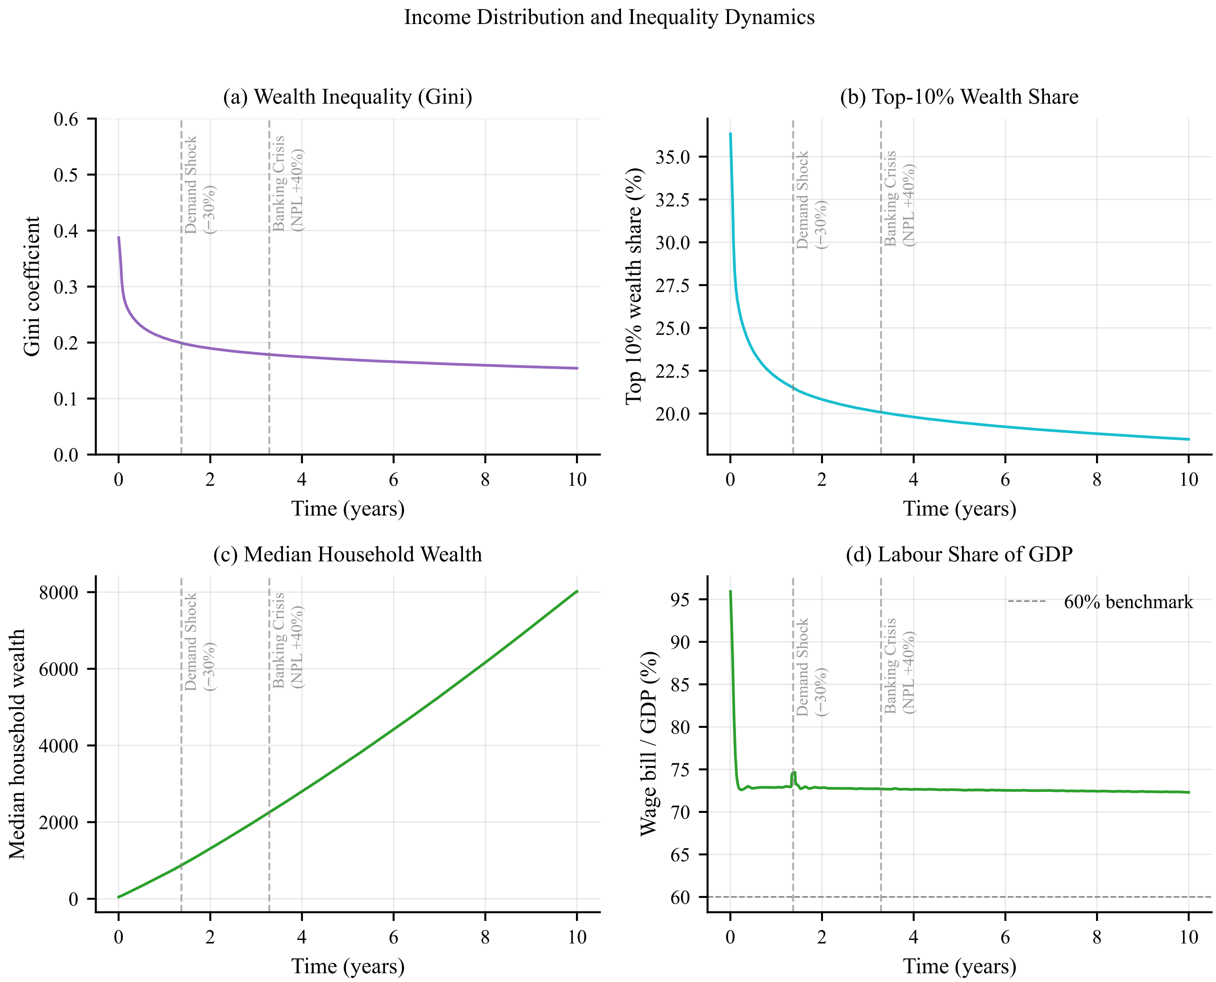

Figure A.5: Inequality dynamics. Panel (a): Gini coefficient of household wealth. Panel (b): top-10 % wealth share. Panel (c): median household wealth. Panel (d): wage share of nominal GDP. All series show 30-day rolling averages.

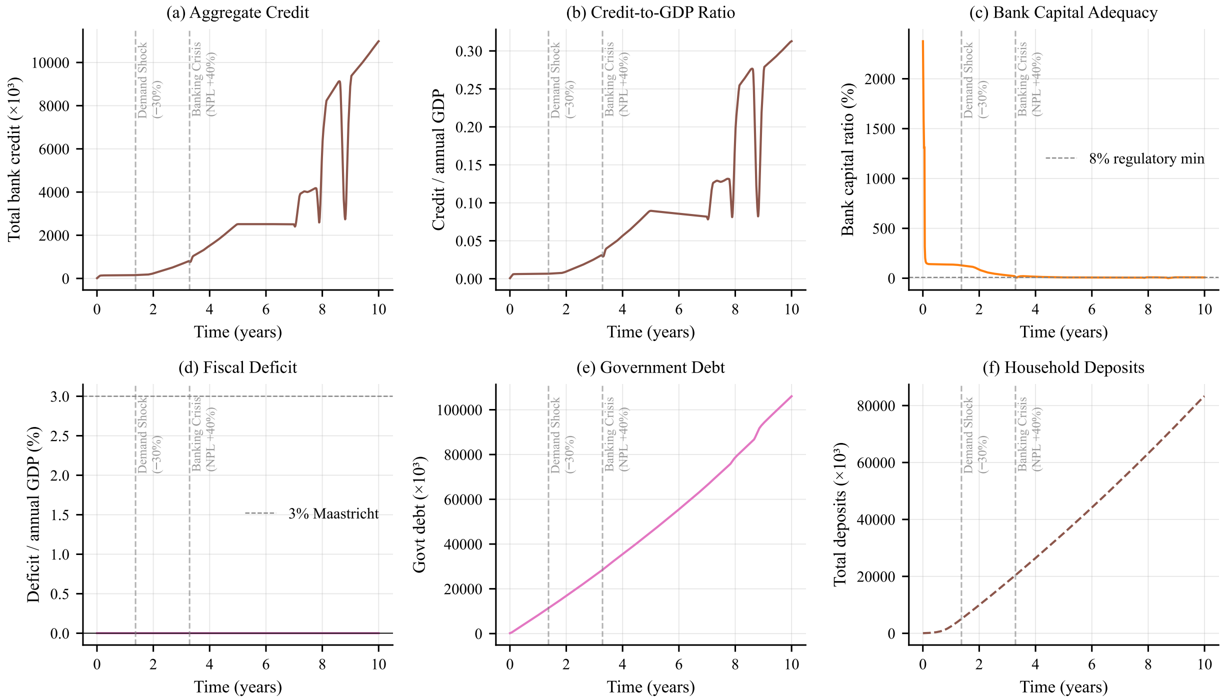

Figure A.6: Financial and fiscal dynamics. From top left: total corporate credit, credit-to-GDP ratio, bank capital ratio, government deficit-to-GDP, cumulative government debt, and household deposits.

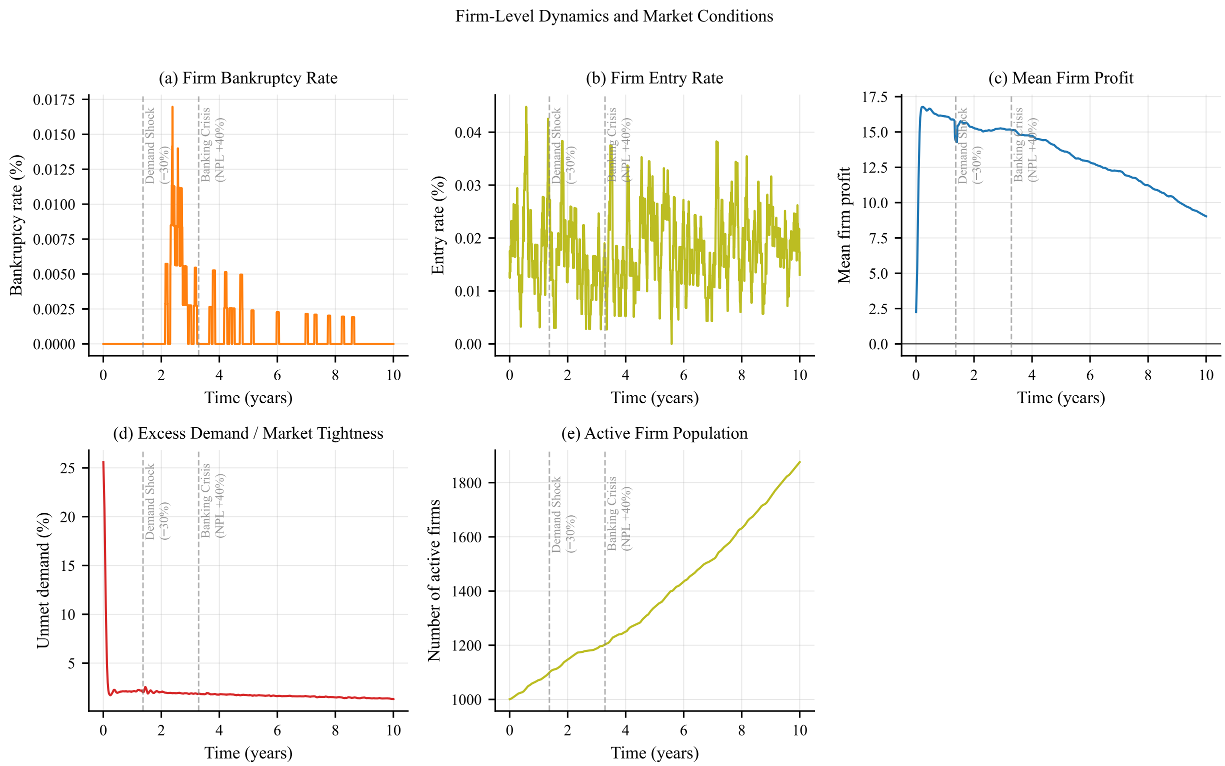

Figure A.7: Firm dynamics: daily firm bankruptcy rate (annualised), entry rate, mean firm profit, and unmet demand fraction. The highest bankruptcy episode occurs in the demand-shock window (years 1.4–3.3); post-crisis bankruptcy remains low under the recalibrated parameters (Section 4.4).

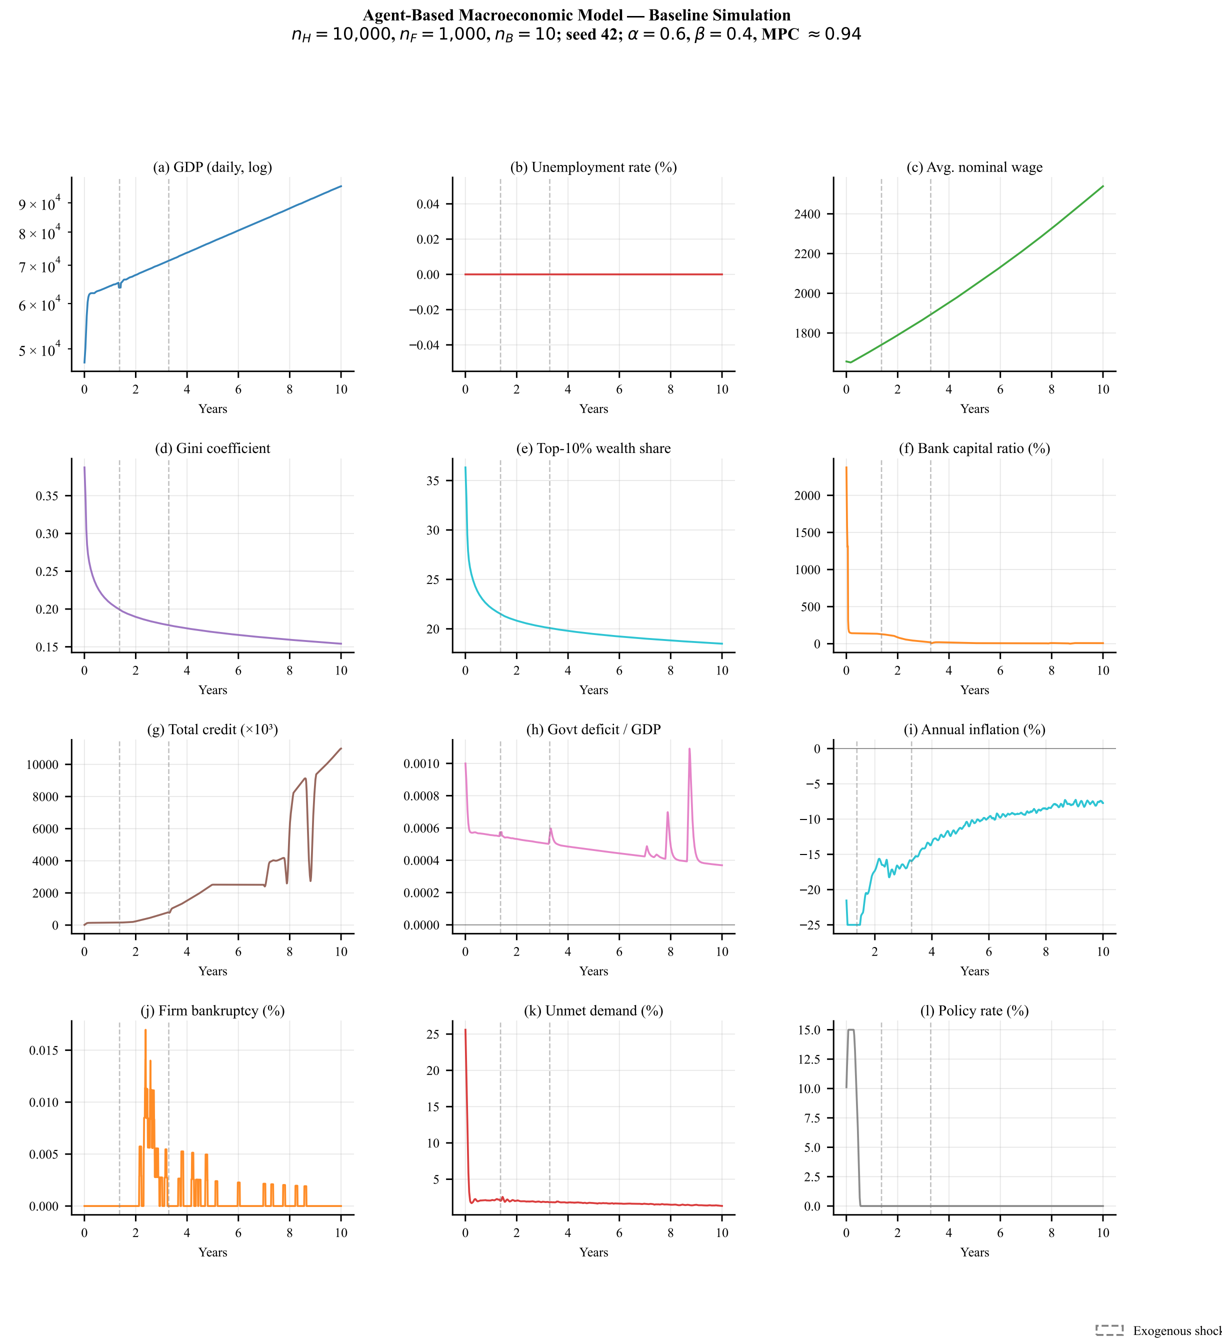

Figure A.8: Four-panel summary: nominal GDP, unemployment rate, Gini coefficient, and firm bankruptcy rate. Layout follows the presentation in Dosi et al. [13].

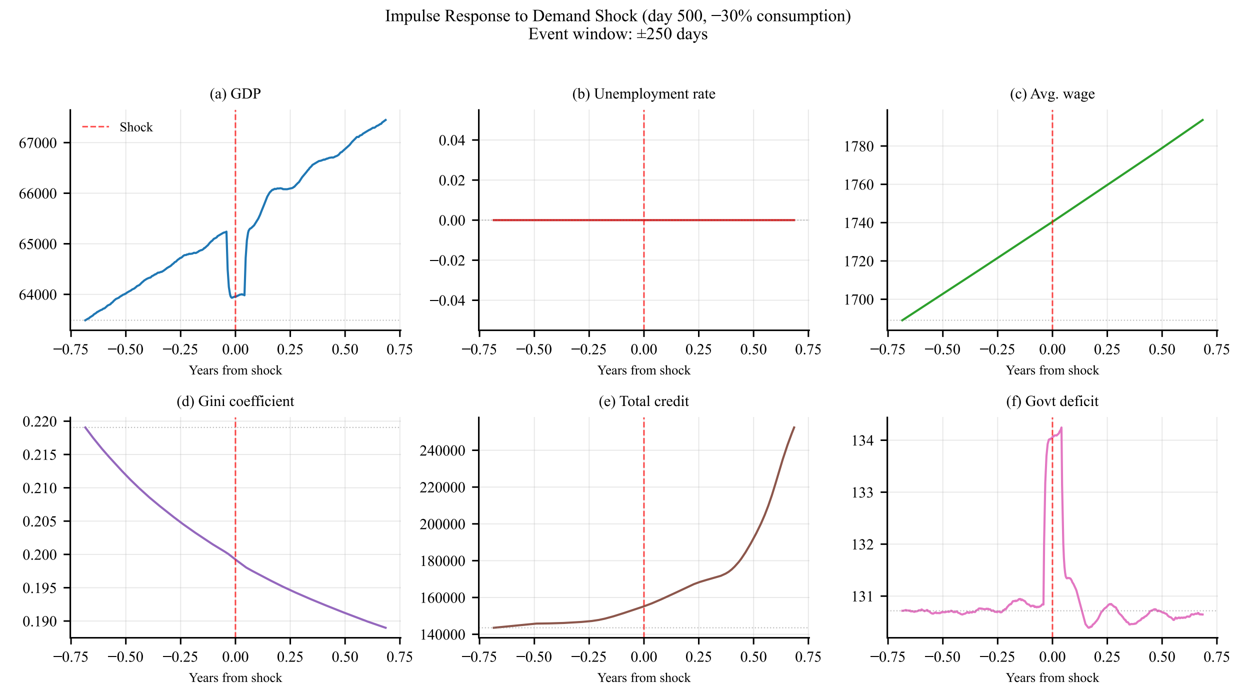

Figure A.9: Impulse responses to the demand shock at day 500 (±250-day window). Each panel shows a key macro variable centred on the shock date. The red dashed line marks the shock day; the dotted horizontal line the pre-shock mean.

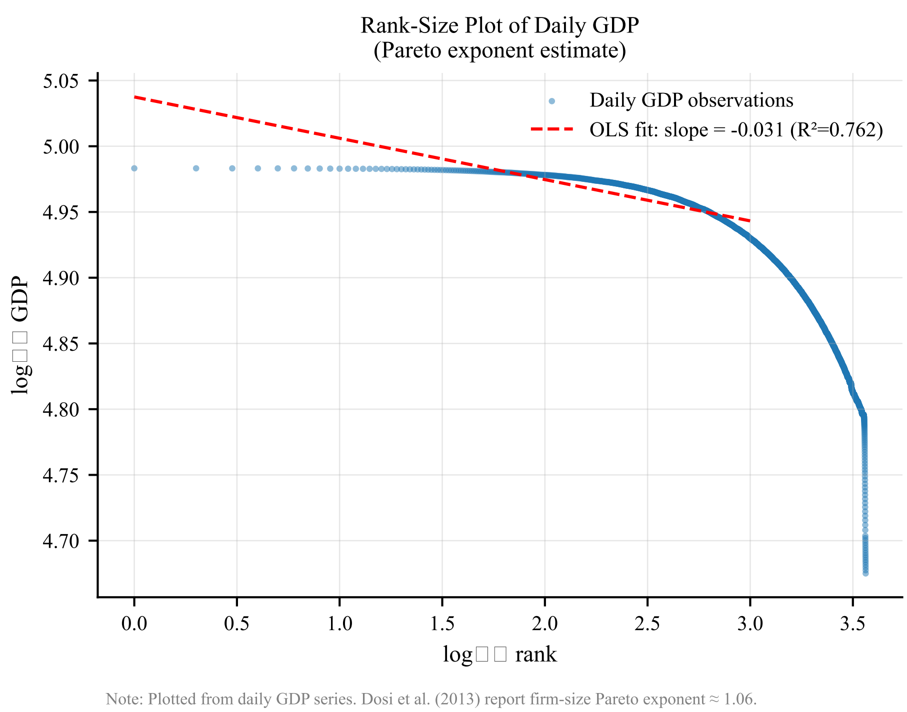

Figure A.10: Log–log rank-size plot of daily GDP observations used as a proxy for the firm-size distribution. The OLS-fitted slope in log–log space estimates the Pareto exponent of the output distribution.

Appendix B. Key Implementation Details

Wage units.. All wages are stored as annualised quantities; daily cash flows use $W_{j,t}/365$. This single convention prevents the unit mismatches that cause spurious 365-fold amplification of wage-drift dynamics (a bug that can cause GDP to collapse within 15 days if the annual drift rate is inadvertently applied to a daily loop).

Budget pre-warming.. At $t = 0$, the consumption budget is set to the steady-state value $c_{i,0} = \tilde{y}_{i,0} \cdot \text{MPC}$ rather than zero. Without this initialisation, the habit term $\eta c_{i,-1} = 0$ suppresses demand for the first 30 days, creating artificial firm losses that can trigger a wave of bankruptcies in the initialisation period.

Social transfers before goods market.. Transfers τ_t are paid at the start of each day and included in the income $y_{i,t}$ that enters the consumption budget formula. If transfers are paid after goods-market clearing, they do not affect today's consumption and the fiscal multiplier is zero by construction.

Entry rate uses alive count.. The Poisson intensity for firm entry uses the count of currently solvent firms N_F^{alive}, not the total length of the firm vector (which includes deceased-but-stored entries). Using the vector length causes exponential growth of the firm population ($> 10^7$ firms after 6 years with the original entry rate of $\nu = 0.005$), consuming > 7 GB of RAM and making completion of the simulation infeasible.

Appendix C. Monte Carlo Robustness

Figure C.11 summarises the distribution of twelve key moments across 50 Halton-seed runs of the baseline configuration. The violin plots show the full empirical distribution; the horizontal lines indicate the 2.5th, 50th, and 97.5th percentiles (95 % CI).

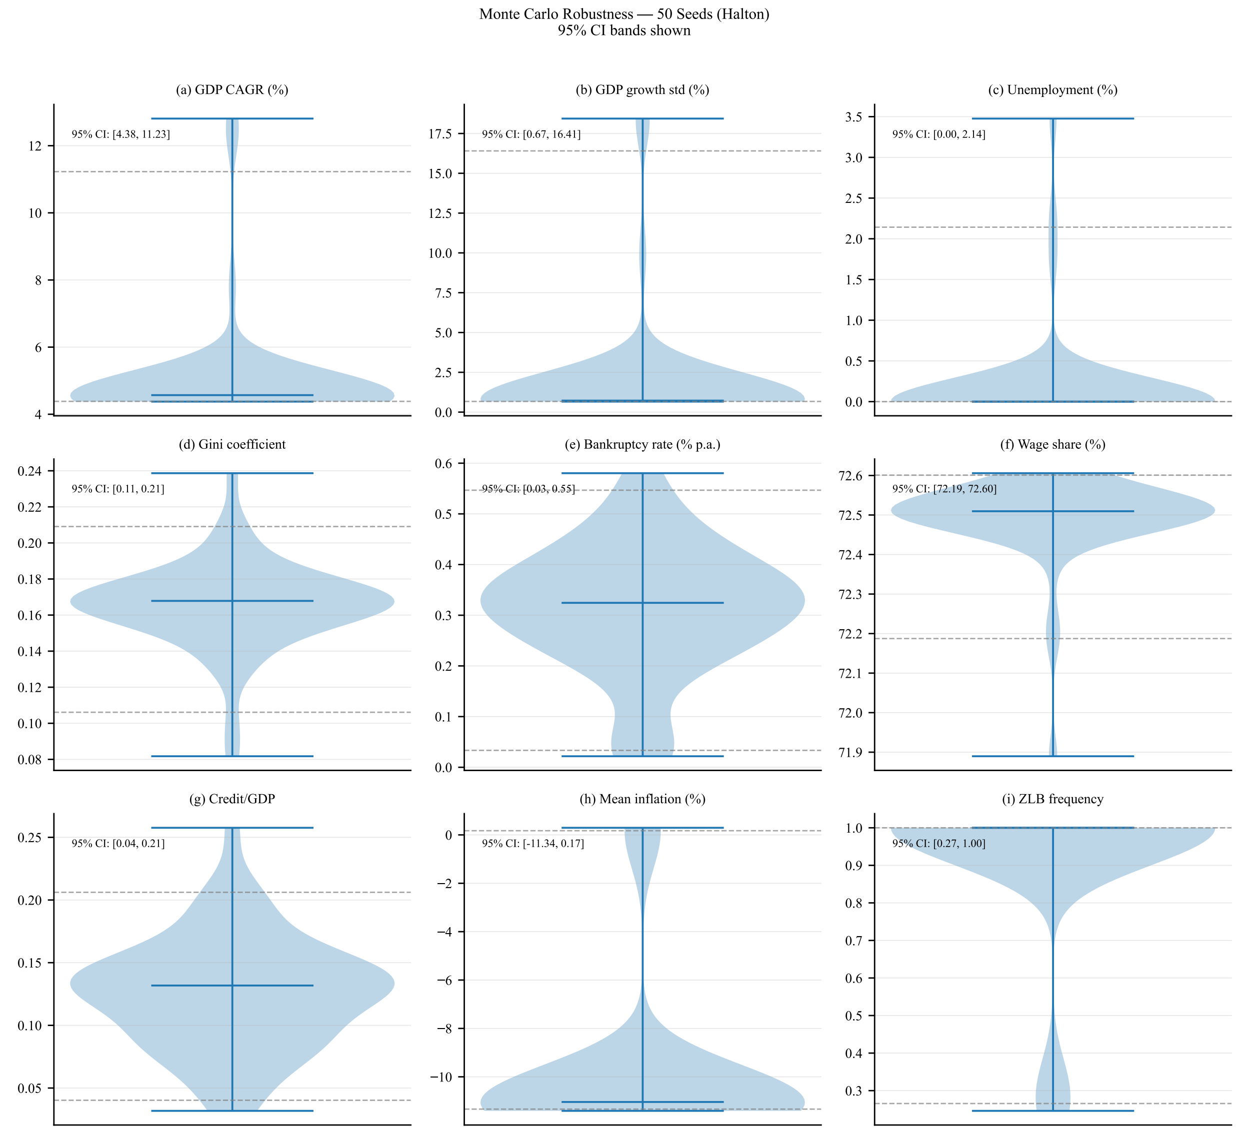

Figure C.11: Monte Carlo robustness: violin plots of key moments across 50 Halton seeds. Rows: business-cycle statistics (top), firm dynamics (middle), distributional moments (bottom). Horizontal lines show 95 % confidence intervals.

Appendix D. Sensitivity Analysis

One-at-a-time (OAT) sensitivity runs vary each of the five free parameters—ξ, θ_f^H, δ_W, χ, κ_P—over a four-point grid with five Halton seeds per grid point (100 runs per parameter). Results from `results/sweeps/` are summarised in companion figures (not shown in this version).

Appendix E. Initial Distribution Experiments

The household wealth and firm-size distributions at $t = 0$ are varied across three functional forms: Pareto (baseline), Uniform, and Log-normal. Two Pareto shape parameters ($\alpha_H \in \{1.2, 1.8\}$ and $\alpha_F \in \{1.0, 1.5\}$) are also swept. All experiments use five Halton seeds. Results from `results/distributions/` indicate that the qualitative behaviour of the model is robust to the choice of initial distribution, while the quantitative inequality moments (Gini, top-10 % share) are sensitive to the Pareto tail index.

Appendix F. Shock Experiments: Impulse-Response Figures

Figures for the supply shock, fiscal stimulus, fiscal austerity, and monetary ZLB scenarios described in Section 4.11 are available from `results/shocks/`. Each figure follows the impulse-response format of Figure A.9, with a ±250-day window centred on the shock date.

Appendix G. Supporting Theoretical Propositions

The following propositions formalise the relationships between the bounded-rational and optimal decision rules discussed in Section 2.7. They are standard results applied to the parameterisation of this model; the proofs follow directly from the cited literature and are included for completeness.

Proposition 4 (Euler–habit approximation gap). *Under CRRA utility with $\gamma = 2$ and a daily deposit rate $r \approx 0$, the Hall Euler rule* (17) *and the habit-income rule coincide to first order when MPC $= 1 - \beta(1 + r)^{1-1/\gamma} \approx 1 - \beta$. The two rules diverge in proportion to the curvature of the marginal utility schedule; at empirically plausible $\gamma \in [1.5, 3]$ the gap in mean consumption is bounded above by $\frac{1}{2}\gamma\,\mathrm{Var}(c)/\bar{c}^2$.*

Proposition 5 (q-theory vs. profit-rate investment). *The Tobin-q rule and the profit-rate rule generate identical investment rates when the firm discount rate equals the internal rate of return on capital. The two rules differ whenever market value deviates from replacement cost;*

in the single-sector ABM firm value is unobservable, so the bounded-rational rule provides an empirically tractable substitute.

Proposition 6 (ZLB under optimal monetary policy)**.** *Under the Woodford optimal Taylor rule with* $\psi_\pi^* \approx 1 + \lambda_y/(\lambda_\pi \kappa^2)$*, the model economy still converges to the ZLB whenever supply-side deflation dominates: the higher inflation coefficient amplifies the policy response but cannot prevent the ZLB if the equilibrium real rate is negative. The Eggertsson–Woodford commitment policy reduces the expected ZLB episode duration by* $T^* = \ln(\pi^*/\pi)/[\kappa(\psi_\pi - 1)]$ *days.*

Proposition 7 (Stiglitz–Weiss credit rationing)**.** *The MLE-calibrated Stiglitz–Weiss loan rate* $r^* = r_{safe} + 1/\hat{\lambda}$ *lies strictly above the bounded-rational lending rate whenever* $\hat{\lambda} < (r_{lending} - r_{safe})^{-1}$*. Because the ABM operates near full employment—suppressing default risk—the MLE hazard is large and the Stiglitz–Weiss correction is small; credit rationing is therefore not the primary driver of low credit-to-GDP in the baseline.*