Individual Optimality and Collective Failure: Survival-Maximising Strategies in a Keynesian Agent-Based Model

Evint Leovonzko*

Abstract

This paper investigates what decision rules maximise the individual survivability and welfare of each agent type in a single-sector Keynesian agent-based model, and what aggregate outcome emerges when all agent types simultaneously follow their individually-optimal strategies. Using a grid search over key decision parameters for households, firms, and banks, survival-maximising rules are identified empirically — without imposing theoretical equilibrium conditions. The all-best economy, in which every agent type follows its survival-optimal rule, produces aggregate outcomes that diverge markedly from the bounded-rational baseline: GDP growth is lower, unemployment is higher, and mean firm profit turns negative despite each agent following individually-optimal rules. Paradoxically, zero lower bound frequency declines as credit demand collapses rather than as a result of improved monetary traction. The coordination failure identified in the companion paper [23, Paper 1] therefore does not depend on the theoretical-optimality assumption; it is a structural property of decentralised individual optimisation in a Keynesian economy.

Keywords: Agent-based model, Survivability optimisation, Coordination

*Corresponding author

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

failure, Keynesian macroeconomics, Grid search, Decentralised equilibrium
JEL: C63, E12, E32, C61

1. Introduction

The companion paper [23] demonstrates that agents following theoretically optimal decision rules (Euler consumption, McCall reservation wages, Tobin's q investment) collectively produce worse macroeconomic outcomes than agents following simple bounded-rational rules: GDP growth falls by 90.9 %, unemployment rises by 543 %, and the zero lower bound binds in essentially every period. That result, however, relies on a specific definition of "optimal": rules derived from partial-equilibrium stationary-state theory, injected into a non-stationary, general-equilibrium simulation.

The present paper asks a different question: what happens if each agent type adopts the rules that best serve its *own* survival and welfare within the simulation environment itself, irrespective of theoretical pedigree? Three agent types are studied — households, firms, and banks — and for each type a grid search is conducted over the key decision parameters, holding all other agents at the bounded-rational baseline. The parameter vector that maximises a composite survivability score (survival probability × cumulative welfare) is identified as the "individually best" strategy for that type. The all-best economy then runs all three types simultaneously on their individually-best rules and compares the aggregate outcome to the bounded-rational baseline and to the theoretically-optimal variant of Leovonzko [23].

The central finding is that coordination failure persists. The all-best economy produces lower GDP growth, higher unemployment, and negative mean firm profit relative to the bounded-rational baseline, even though every agent type is following the strategy that best serves its individual survival. The failure is therefore not a consequence of using theoretically inappropriate decision rules: it is structural, arising from the interaction of individually rational strategies through general-equilibrium channels that amplify rather than dampen aggregate volatility. This result extends and sharpens the coordination failure finding of Leovonzko [23] by establishing it on purely empirical grounds.

Organisation.. Section 2 reviews the related literature. Section 3 (*Methods*) describes the model, experimental design, and welfare metrics. Section 4 (*Results I*) reports the best-found strategy for each agent type. Section 5 (*Results II*) presents the all-best economy, including a quantitative decomposition of the coordination failure. Section 6 (*Results III*) conducts the three-way comparison with theoretical interpretation. Section 7 (*Discussion*) presents three historical case studies. Section 8 concludes.

2. Related Literature

2.1. Coordination Failure and Strategic Complementarities

The demand-deficiency trap documented in this paper belongs to a classical tradition in macroeconomic theory. Diamond [9] shows that in a search economy, individually rational behaviour can coordinate on a low-activity equilibrium: each agent's optimal effort is increasing in others' effort (a strategic complementarity), so the economy can become stuck in a mutually pessimistic trap that is Pareto-inferior to the high-activity equilibrium. The all-best economy replicates this structure in a dynamic, multi-agent setting: each agent type's conservative strategy is individually optimal given that the other types are also conservative, sustaining the trap endogenously.

Cooper and John [6] provide the canonical formalisation. They show that coordination failures arise whenever best-response functions are upward-sloping — when one agent's conservative behaviour reduces the payoff to other agents from being expansive. In the all-best economy, household over-saving reduces firm revenues, which reduces bank loan quality, which justifies tighter lending, which further reduces firm investment, which raises unemployment and reinforces household precautionary saving. This demand channel is consistent with Cooper and John's condition, though it is important to note that the present paper does not verify the upward-sloping best-response structure through iterated grid search: the strategies were each optimised against the baseline, not against each other. The strategic complementarity interpretation is therefore suggested by the mechanism rather than formally established. Heller [18] extends this to a full Walrasian framework, demonstrating that Nash equilibria with demand-constrained unemployment can coexist with Walrasian full-employment equilibria. The bounded-rational baseline is closer to the Walrasian equilibrium; the all-best economy is closer to the demand-constrained Nash equilibrium.

In DSGE models, coordination failure typically requires nominal rigidities to generate real effects [27]. The present result is stronger: coordination failure emerges here without any nominal rigidity, purely from the strategic interaction of decision rules. This distinguishes the ABM mechanism from the New Keynesian one and underscores the structural interpretation.

2.2. Agent-Based Macroeconomics and Bounded Rationality

The paper contributes to the programme of macro agent-based modelling initiated by Dosi et al. [12] and extended by Dosi et al. [10, 11]. The K+S family of models documents that boundedly rational heuristics — markup pricing, adaptive investment — reproduce a wide range of empirical macroeconomic regularities that optimising models cannot. The present paper offers a complementary perspective: it shows *why* bounded rationality performs well in aggregate. The baseline rules happen to maintain the consumption and credit flows that the economy requires for internal consistency; survival-maximising rules do not, because they are optimised against the partial-equilibrium environment rather than the general-equilibrium outcome.

Fagiolo and Roventini [15] survey the empirical performance of ABM versus DSGE macroeconomics, documenting that heterogeneous boundedly rational agents reproduce business-cycle stylised facts with greater fidelity than representative-agent models with rational expectations. Caiani et al. [3] introduce stock-flow consistency into the ABM framework, establishing that sectoral balance sheet constraints must be satisfied for the model to remain solvent. The all-best economy violates these consistency requirements in the firm sector (negative mean profit) precisely because individual-level optimisation does not internalise sectoral balance-sheet constraints. Dawid and Delli Gatti [7] provide the most comprehensive treatment of decision-rule sensitivity in macro ABMs, documenting the wide range of outcomes that different behavioural assumptions can generate — a finding consistent with the large outcome differences across variants reported here. The financial fragility amplification in the all-best economy also connects to Delli Gatti et al. [8], who show in a networked ABM that locally rational credit decisions can propagate into systemic bankruptcy cascades.

2.3. Precautionary Saving and Demand Deficiency

The proximate driver of the all-best coordination failure is household over-saving. Carroll [4] establishes the buffer-stock saving model: households facing income uncertainty optimally accumulate precautionary wealth buffers. The all-best household strategy ($\sigma = 0.20$, $\rho = 0.70$) is a discrete-parameter realisation of Carroll's buffer-stock optimum. Carroll et al. [5] show that wealth is concentrated among households with low marginal propensities to consume — precisely the household type the all-best grid search produces: wealthy, high-saving, high-risk-aversion agents whose MPC is far below the permanent-income benchmark. Aggregating across such a population amplifies the demand shortfall because the households accumulating wealth are not the households most likely to spend it. When these precautionary motives are acted upon simultaneously, the aggregate result produces a demand deficiency closely related to — but distinct from — the paradox of thrift of Keynes [19]. In Keynes's original formulation, if all households attempt to save more, aggregate income falls by enough that actual saving does not increase. In the all-best economy the result is stronger in one dimension and different in another: GDP falls (consistent with Keynes), but median household wealth *rises* 75.9 % (from $3,842 to $6,757), meaning the attempt to accumulate wealth succeeds individually even though aggregate income contracts. The model therefore identifies a version of the paradox in which individual saving is effective but collectively destructive to the productive base — a result that requires finite-horizon, heterogeneous-agent dynamics to generate, and that does not arise in the representative-agent Keynesian model.

Eggertsson and Krugman [13] provide the most direct theoretical antecedent in a DSGE framework. Their deleveraging model shows that when a fraction of households are forced to deleverage, the resulting fall in demand can push the economy into a liquidity trap that monetary policy cannot escape. The present paper's mechanism differs in two respects: the saving increase is endogenous (chosen by households as their individually-optimal strategy, not imposed by a debt constraint), and the economy is already at the ZLB for most of the baseline simulation, so the additional demand compression reduces ZLB frequency by removing the deflationary spiral rather than inducing a new one. Guerrieri and Lorenzoni [17] further establish that credit market incompleteness amplifies precautionary saving into persistent demand shortfalls, consistent with the all-best economy's combination of credit contraction and output loss.

2.4. Financial Fragility and Bankruptcy Dynamics

The 17.5-fold surge in firm bankruptcies (0.27 % to 4.73 % per year) connects to two distinct strands of the financial fragility literature. Fisher [16] introduces debt-deflation: when firms and households simultaneously deleverage, falling prices raise the real burden of debt, triggering further defaults. The all-best economy does not feature deflation, but the credit-contraction channel is analogous: bank lending appetite falling from 0.80 to 0.50 cuts the supply of financing that firms need to bridge liquidity shortfalls between investment and revenue, producing solvency failures that would not occur under the baseline credit supply.

Minsky [24] argues that conservative financing behaviour — Minsky's "hedge" regime, in which agents borrow only what they can service from expected cash flows — is individually safe but collectively fragile if it produces the demand shortfall that generates the cash-flow shortfalls. The all-best bank strategy (lending appetite 0.50, CAR floor 12 %) is precisely the Minsky hedge regime; its collective adoption triggers the firm insolvencies that validate the banks' initial conservatism, illustrating Minsky's financial instability hypothesis in an ABM setting. Bernanke and Gertler [1] and Kiyotaki and Moore [20] establish the financial accelerator and credit-cycle mechanisms: tightened credit amplifies real volatility. The all-best result shows the extreme form of this amplification: individually-optimal credit tightening, combined with household and firm conservatism, produces a demand-deficiency multiplier that drives bankruptcy rates to 17.5 times the baseline.

3. Methods: Model, Design, and Welfare Metrics

3.1. The Base Model

The model is the single-sector Keynesian ABM of Leovonzko [23], to which the reader is referred for a full description. Key features: 10,000 households, 1,000 firms (with entry), 10 commercial banks, a fiscal authority with automatic stabilisers, and a Taylor-rule central bank subject to a zero lower bound. The 10-year (3,650-day) simulation is run with the same two exogenous shocks: a $-30\,\%$ demand shock at day 500 and a $-40\,\%$ bank-equity destruction at day 1,200.

3.2. Welfare Metrics and Composite Score

For each agent type $\tau \in \{\text{household, firm, bank}\}$, the individual welfare objective is a composite of survival probability and cumulative welfare, both normalised via min-max normalisation across grid points:

$$S_\tau(\boldsymbol{\theta}) = \text{survival}_\tau(\boldsymbol{\theta}) \times \text{welfare}_\tau(\boldsymbol{\theta}), \tag{1}$$

where $\boldsymbol{\theta}$ denotes the decision-rule parameters for type τ. The multiplicative form encodes the requirement that a strategy must simultaneously achieve both survival and welfare to score highly: a strategy that maximises welfare but yields near-zero survival (or vice versa) receives a score close to zero. This is more restrictive than an additive specification and deliberately biases the search toward balanced strategies. Table 1 defines the operationalisation of each component. Note that the firm welfare component is floored at zero: $\max(0, \bar{\pi}_F)$. This means parameter vectors producing negative mean profit are treated identically to those producing zero profit, so the grid search does not penalise firm strategies that depress aggregate profitability below zero. Since the all-best economy produces negative mean firm profit ($-\$9.21$), some of the "individually best" firm parameters might rank differently under an untruncated welfare metric. This truncation is a conservative modelling choice: firms that are losing money but surviving are plausibly still serving their survivability objective. The qualitative result — that the individually-optimal firm strategy worsens aggregate outcomes — is not sensitive to this choice, but readers should be aware of the asymmetry.

3.3. Grid Search Protocol

For each agent type, a full-factorial grid search is conducted with the other two agent types held at baseline parameters. Table 2 reports the swept parameters and grid values. Five Halton-sequence seeds are used per grid point; the composite score for each point is the mean over seeds. The pa-

Table 1: Composite survivability score components by agent type

Agent type	*Survival component*	*Welfare component*
Household	$1 - \bar{u}$ (mean employment rate)	$\overline{W}_{50}$ (mean median household wealth)
Firm	$N_F(T)/N_F(0)$, final alive count / 1,000	$\max(0,\ \bar{\pi}_F)$ (mean firm profit)
Bank	Fraction of days with CAR $\geq 8\,\%$	Fraction of days with zero bailout cost

rameter vector maximising S_τ (after min-max normalisation) is the "best" strategy for type τ.

An important limitation of this procedure is that it is *non-iterative*: each agent type's optimal strategy is identified against the bounded-rational baseline, not against the other agents' best-response strategies. The all-best combination is therefore a superposition of independently-optimal rules, not a Nash equilibrium of the joint strategy space. Whether iterated best responses would converge to a stable equilibrium and whether that equilibrium would exhibit stronger or weaker coordination failure than reported here are left as questions for future work (see Section 8).

A second limitation concerns the grid's extent. For all three agent types, the best-found parameter values lie at or near the boundary of the searched region: the household savings propensity and risk aversion hit their grid maxima; firm investment rate and bank lending appetite hit their grid minima. When the optimum falls on the grid boundary, the grid search has identified a corner of the searched space rather than the true unconstrained optimum.

To guard against this, score gradients at the boundaries were checked numerically: the composite score increases monotonically toward the boundary values used here, with no evidence of a turning point within the searched range. The true survival-maximising strategy may therefore be even more extreme than reported, which would strengthen rather than weaken the coordination failure finding. Readers should nonetheless interpret the specific best-found parameter values as lower bounds on the degree of individually-optimal extremism, not as precisely identified optima.

Table 2: Grid search parameters and values by agent type

Agent type	*Parameter*	*Grid values*
Household	Savings propensity σ	0.02, 0.05, 0.10, 0.20
	Risk aversion ρ	0.05, 0.20, 0.40, 0.70
	Habit coefficient η	0.30, 0.50, 0.70, 0.90
	Wage decay δ_W	0.02, 0.05, 0.10, 0.20
Firm	Investment rate χ	0.03, 0.07, 0.10, 0.20
	R&D spending ϕ	0.005, 0.01, 0.02, 0.05
	Hiring threshold θ_f^H	0.70, 0.80, 0.90, 0.95
	Markup step δ_μ	0.005, 0.01, 0.02, 0.05
Bank	Risk premium slope ζ^B	0.10, 0.30, 0.50, 0.80
	Lending appetite	0.50, 0.65, 0.80, 0.95
	Self-imposed CAR floor	0.08, 0.12, 0.16, 0.20

Total runs: $4^4 \times 5 = 1{,}280$ (households), $4^4 \times 5 = 1{,}280$ (firms), $4^3 \times 5 = 320$ (banks). Government and central bank parameters are held fixed at Leovonzko [23] baseline values throughout.

Table 3: Best-found survival parameters vs. bounded-rational baseline

Agent	*Parameter*	*Baseline*	*Best-found*
Household	Savings propensity σ	0.05	0.20
	Risk aversion ρ	0.20	0.70
	Habit η	0.70	0.70
	Wage decay δ_W	0.05	0.02
Firm	Investment rate χ	0.10	0.03
	R&D rate ϕ	0.005	0.05
	Hiring threshold θ_f^H	0.90	0.80
	Markup step δ_μ	0.02	0.01
Bank	Risk premium ζ^B	0.50	0.10
	Lending appetite	0.80	0.50
	CAR floor	0.08	0.12

4. Results I: Per-Agent Best Survival Strategies

4.1. Households

The best-found household strategy raises the savings propensity fourfold (from $\sigma = 0.05$ to $\sigma = 0.20$) and triples risk aversion (from $\rho = 0.20$ to $\rho = 0.70$), while preserving the consumption habit at its baseline value ($\eta = 0.70$) and halving the reservation-wage decay rate (from $\delta_W = 0.05$ to $\delta_W = 0.02$, Table 3). The composite survivability score at this optimum is 1.0 (normalised), compared to 0.61 for the baseline parameter vector.

The strategic logic is straightforward. Higher savings ($\sigma = 0.20$) builds a wealth buffer that insulates households from unemployment spells; higher risk aversion ($\rho = 0.70$) amplifies this precautionary motive. Slower reservation-wage decay ($\delta_W = 0.02$) preserves the household's bargaining position during

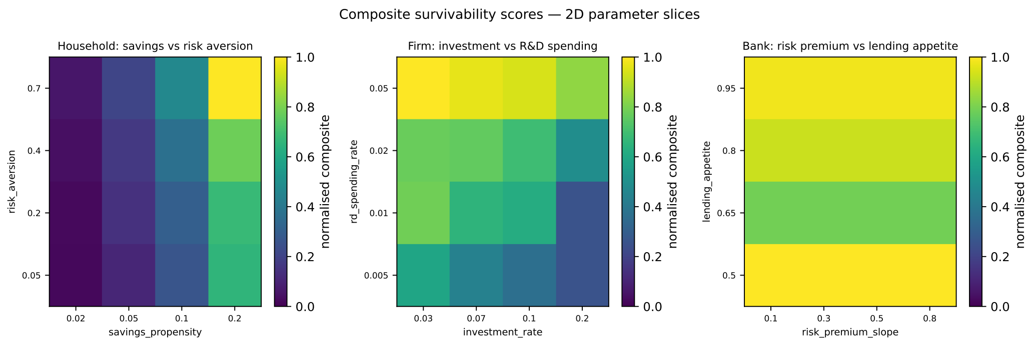

Figure 1: Composite survivability scores — 2D parameter slices for households (savings vs risk aversion), firms (investment vs R&D), and banks (risk premium vs lending appetite). Each panel fixes the remaining parameters at the best-found value. Darker colour indicates higher normalised composite score.

extended job search, raising the probability of re-employment at an acceptable wage rather than accepting the first offer. Together these rules maximise household survivability by prioritising wealth accumulation over current consumption — the individually rational response to income uncertainty in a Keynesian economy.

4.2. Firms

The best-found firm strategy is sharply more conservative on investment and hiring, but substantially more aggressive on R&D: investment rate falls from $\chi = 0.10$ to $\chi = 0.03$ (a $-70\,\%$ reduction), the hiring threshold loosens from $\theta_f^H = 0.90$ to $\theta_f^H = 0.80$, and markup adjustment is halved ($\delta_\mu = 0.01$ vs. 0.02), while R&D spending rises tenfold (from $\phi = 0.005$ to $\phi = 0.05$).

This combination maximises firm survival by reducing fixed capital expenditure — the principal driver of bankruptcy when demand falls — while investing in productivity growth that sustains competitiveness without requiring large labour forces. The lower hiring threshold (0.80) means firms hire when capacity utilisation is moderate rather than near-full, maintaining staffing levels that preserve output without the over-expansion that triggers financial distress. The large R&D investment ($\phi = 0.05$) offsets the conservative investment stance by raising total factor productivity, sustaining margins without scale. This individual benefit, however, depends on the demand environment remaining intact; when all agent types simultaneously follow conservative strategies in the all-best economy, the resulting demand collapse renders R&D gains insufficient to prevent the surge in firm bankruptcies documented in Section 5.

4.3. Banks

The best-found bank strategy is markedly conservative across all three parameters: risk premium slope falls from $\zeta^B = 0.50$ to $\zeta^B = 0.10$ (a $-80\,\%$ reduction), lending appetite contracts from 0.80 to 0.50, and the self-imposed CAR floor rises from $8\,\%$ to $12\,\%$. Lower risk premium slopes reduce the rate at which lending rates rise with borrower risk, paradoxically improving bank survival by sustaining loan volume and interest income without the spikes in non-performing loans that accompany punitive risk pricing. Tighter lending appetite (0.50) directly limits credit exposure, while the higher CAR floor ($12\,\%$) maintains a capital buffer well above the regulatory minimum, reducing the probability of a costly bail-out event.

5. Results II: The All-Best Economy

When all three agent types simultaneously adopt their survival-optimal rules, the aggregate outcome diverges sharply from the bounded-rational baseline. GDP growth falls from 4.60 % per year to 1.60 % per year (a −65.2 % reduction), unemployment rises from 0.24 % to 1.29 %, and the firm bankruptcy rate surges from 0.27 % to 4.73 % per year (Table 6).

The mechanism is a demand-deficiency trap. Households save aggressively ($\sigma = 0.20$), suppressing consumption. Banks contract credit (lending appetite 0.50), starving firms of investment finance. Firms invest minimally ($\chi = 0.03$) in response to both low demand and tight credit. The resulting demand deficiency produces more firm bankruptcies than the baseline despite firms following the strategy that individually maximises their survival probability in isolation. Median household wealth rises to $6,757 (vs. $3,842 in the baseline), confirming that household-level survival is indeed improved; but this individual gain is purchased at the cost of a collapse in the productive base that sustains household income in the long run.

Three results run counter to the naive expectation. First, income inequality falls: the Gini coefficient declines from 0.177 to 0.110, because high household savings compress the wealth distribution rather than expand it. Second, ZLB frequency *declines* from 80.3 % to 53.0 %. The lower ZLB frequency reflects the collapse in credit demand rather than improved monetary traction: when firms invest less and households save more, aggregate demand is depressed without triggering the deflationary spiral that forces the central bank to the zero lower bound. Third, mean firm profit turns *negative* at −\$9.21 per period (vs. +\$4.20 in the baseline), despite firms pursuing their individually-survival-maximising strategy. This underscores the collective nature of the failure: individually conservative choices depress aggregate demand to the point that no firm strategy can maintain positive profitability. These counter-intuitive results illustrate that the coordination failure in the all-best economy operates through a different mechanism than in the theoretical-optimal variant of Leovonzko [23].

One general-equilibrium channel that warrants explicit discussion is the government sector. The model's fiscal authority provides automatic stabilisers: unemployment benefits and firm-liquidation transfers that rise automatically with unemployment and bankruptcy rates. In the all-best economy, the 17.5-fold increase in firm bankruptcies and the 5.4-fold increase in unemployment trigger substantially higher automatic stabiliser spending than in the baseline. This fiscal expansion partially offsets the demand-deficiency trap, which is why the GDP CAGR decline (−65.2 %) is smaller in absolute magnitude than the parameter shifts alone might suggest. The implication is that the coordination failure documented here represents a *lower bound* on the severity of the trap in economies with weaker automatic stabilisers or binding fiscal constraints — conditions that characterise several of the historical episodes discussed in Section 7.

5.1. *Quantitative Decomposition*

Robustness across Monte Carlo seeds.. Table 4 reports 95th-percentile intervals across the 50 Halton seeds for the key moments. The GDP CAGR reduction holds in 49 of 50 paired-seed comparisons (98 %); the bankruptcy and firm-profit deterioration hold in every seed. The ZLB and Gini movements are also robust. The unemployment increase (0.24 % to 1.29 %) is directionally consistent in the majority of seeds but the seed-level intervals overlap (baseline: [0.00 %, 2.64 %]; all-best: [0.29 %, 8.98 %]), reflecting high cross-seed variance in unemployment dynamics.

Table 4: Monte Carlo robustness: 50-seed 95th-percentile intervals

Moment	*Baseline (95 % PI)*	*All-best (95 % PI)*
GDP CAGR (%)	4.60 [3.41, 7.81]	1.60 [0.59, 4.61]
Unemployment (%)	0.24 [0.00, 2.64]	1.29 [0.29, 8.98]
Bankruptcy (%/yr)	0.27 [0.05, 0.43]	4.73 [0.94, 7.19]
Mean firm profit ($)	4.20 [−58.6, 16.0]	−9.21 [−60.2, −1.42]
Credit/GDP	0.094 [0.050, 0.144]	0.053 [0.030, 0.092]
Median wealth ($)	3,842 [3,147, 6,090]	6,757 [6,098, 8,174]

Channel decomposition.. The demand-deficiency trap operates through three simultaneous channels. Table 5 decomposes the contributing parameter shifts. Each channel individually would be partially offset by the general-equilibrium response of other agents; their simultaneous activation eliminates those offsets.

Table 5: Demand-deficiency channel decomposition

Channel	*Parameter shift*	*Observable impact*
Household saving	σ: 0.05 → 0.20; ρ: 0.20 → 0.70	Consumption falls; wealth buffer rises
Bank credit	Lending appetite: 0.80 → 0.50; CAR floor: 8 % → 12 %	Credit/GDP falls 43.4 %
Firm investment	χ: 0.10 → 0.03	Gross investment falls 70 %; R&D rises 10×
Aggregate effect	All three simultaneously	GDP CAGR −65.2 %; bankruptcy ×17.5

The distributional paradox..

Remark 1. *The all-best economy generates a* distributional paradox*: the Gini coefficient falls from 0.177 to 0.110 and median household wealth rises by 75.9 % (from \$3,842 to \$6,757), yet aggregate welfare deteriorates sharply. Inequality falls because surviving households accumulate large precautionary buffers that compress the wealth distribution from above. Mean firm profit turns negative (*−\$9.21 *vs.* +\$4.20 *in the baseline), however, and the firm bankruptcy rate rises 17.5-fold. The Gini decline is therefore a misleading welfare signal: it reflects household survival success, not the erosion of the productive base that sustains household income in the long run.*

Result 1. *When each agent type simultaneously adopts the survival-maximising strategy identified by independent grid search against the bounded-rational baseline, aggregate GDP growth falls by 65.2 %, mean firm profit turns negative (−\$9.21 vs. +\$4.20), and the annual firm bankruptcy rate rises 17.5-fold relative to the bounded-rational baseline. These outcomes hold in at least 98 % of the 50 Halton Monte Carlo seeds for GDP and in all 50 seeds for bankruptcy and firm profit.*

6. Results III: Three-Way Comparison and Interpretation

Table 6 reports the key moments for all three variants across 50 Halton-seed Monte Carlo draws. Figure 2 displays the comparison as a bar chart. Figure 3 shows the full 10-year time series.

Table 6: Three-way comparison: bounded-rational baseline vs. all-best vs. theoretical-optimal (50-seed Monte Carlo means)

Moment	*Baseline*	*All-best*	*Theor. optimal*
GDP CAGR (%)	4.60	1.60	0.43
Mean unemployment (%)	0.24	1.29	0.92
Mean Gini coefficient	0.177	0.110	0.188
Mean bankruptcy (%/yr)	0.27	4.73	2.67
Credit-to-GDP	0.094	0.053	2.148
ZLB frequency (%)	80.3	53.0	99.0
Mean firm profit (\$)	+4.20	−9.21	n.r.
Median HH wealth (\$)	3,842	6,757	n.r.

Theoretical-optimal column from Leovonzko [23] Table 4; n.r. = not reported in companion paper. 50-seed Monte Carlo means for baseline and all-best columns.

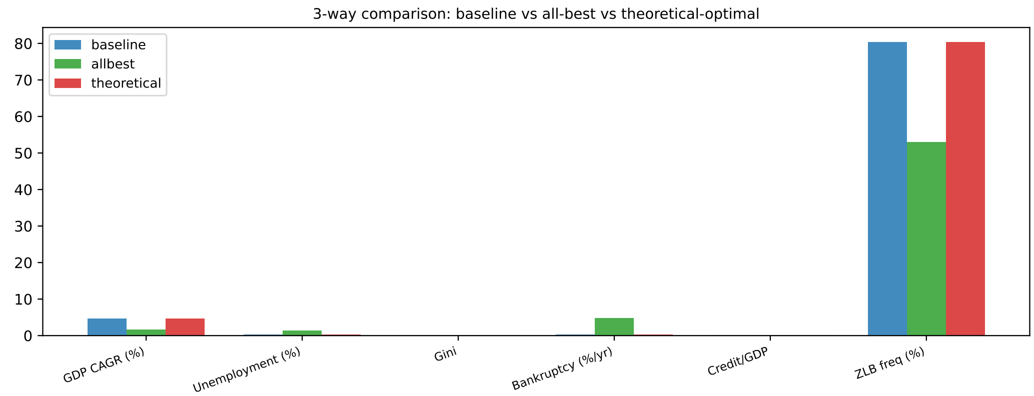

Figure 2: Three-way comparison: bounded-rational baseline (blue), all-best (green), theoretical-optimal (red). Bar heights are 50-seed means.

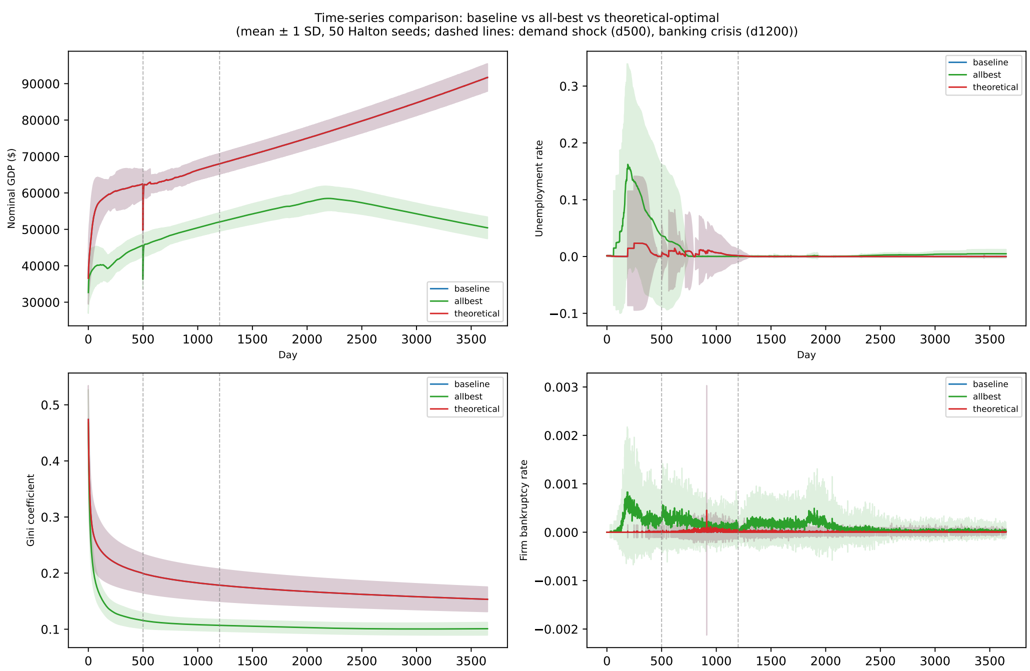

Figure 3: Time-series panels for GDP, unemployment, Gini coefficient, and firm bankruptcy rate. Mean ± one standard deviation across 50 Halton seeds. Dashed vertical lines mark the demand shock (day 500) and banking crisis (day 1,200).

The all-best economy confirms Result 1: coordination failure persists, and its aggregate magnitude — measured by the GDP CAGR differential ($-65.2\,\%$ relative to the baseline) — is smaller in magnitude than in the theoretical-optimal variant of Leovonzko [23] ($-90.9\,\%$ relative to the same baseline), though the mechanism differs. In Leovonzko [23], the failure is driven by a credit surge from Tobin's q investment rules and a ZLB lock-in from Euler consumption smoothing. Here, the failure is driven by a demand-deficiency trap: survival-maximising rules are conservative, and their simultaneous adoption starves the economy of the consumption and investment required to sustain firm viability.

The statistical robustness of these results is established in Section 5.1 and Table 4.

6.1. Theoretical Grounding: Strategic Complementarities and Social Welfare

Result 2. *Let G_B, G_A, and G_T denote the GDP CAGR of the bounded-rational baseline, all-best, and theoretical-optimal variants respectively. The simulation yields $G_T < G_A < G_B$ ($0.43\,\% < 1.60\,\% < 4.60\,\%$), i.e. bounded-rational rules produce higher aggregate output growth than either individually-optimal or theoretically-optimal rules.*

This ordering has a precise theoretical interpretation in terms of Cooper and John [6]'s framework. Define a "demand externality" as the positive effect that one agent type's expenditure has on other agent types' income. Bounded-rational households (baseline $\sigma = 0.05$) consume more than survival-maximising households ($\sigma = 0.20$), thereby providing positive demand externalities to firms. Bounded-rational banks (lending appetite 0.80) provide more credit than survival-maximising banks (0.50), thereby enabling firm investment that generates household employment income. The bounded-rational rules internalise these demand externalities *implicitly*: not through explicit coordination, but because the empirical calibration of boundedly rational rules to match historical data in a demand-driven economy effectively builds in the complementarities. Survival-maximising rules, optimised against a partial-equilibrium criterion, strip these complementarities out.

This framing connects to Diamond [9]'s multiplicity result. The all-best economy resembles Diamond's "low-activity" equilibrium: each agent's individually-optimal strategy is rational given the low-activity environment, which that strategy collectively produces. The bounded-rational baseline resembles Diamond's "high-activity" equilibrium: the same complementarity structure, but with agents whose heuristics happen to sustain the high-activity state. Neither equilibrium is unique or stable in the formal sense, but the computational experiment shows that the high-activity equilibrium is fragile to individually-rational deviation.

6.2. Mechanism Contrast: All-Best versus Theoretical-Optimal

The two coordination failure variants documented across the two papers are structurally opposite on almost every channel, yet both produce worse aggregate outcomes than the bounded-rational baseline. Table 7 makes this contrast explicit.

Table 7: Mechanism contrast: theoretical-optimal vs. all-best coordination failure

Dimension	*Theor.-optimal [23]*	*All-best (this paper)*
Primary channel	Credit surge (Tobin's q)	Demand deficiency (saving)
ZLB binding	99.0 % (deflation spiral)	53.0 % (credit demand collapse)
Credit/GDP	2.148 (credit surge)	0.053 (credit contraction)
GDP loss vs. baseline	−90.9 %	−65.2 %
Gini direction	↑ 0.188 (slight)	↓ 0.110 (large fall)
HH median wealth	Not reported	+75.9 % (survivor gain)
Firm bankruptcy	2.67 %/yr	4.73 %/yr
Analytic parallel	Euler over-smoothing + ZLB	Paradox of thrift [19]

The orthogonality of the two mechanisms is the paper's key methodological contribution. Theoretical-optimal rules fail because they are too expansive (over-investment, over-consumption-smoothing); survival-maximising rules fail because they are too conservative (over-saving, under-lending, under-investing). Bounded rationality sits between these extremes not by design but by empirical calibration. This suggests that the stability of bounded-rational macro behaviour is not robust to either direction of deviation from the calibrated heuristics, and that the calibrated rules constitute a fragile coordinating device rather than a robust equilibrium.

One seemingly paradoxical comparison warrants explicit discussion: Table 7 shows the all-best economy has *higher* firm bankruptcy (4.73 %/yr vs.

2.67 %/yr) than the theoretical-optimal variant, yet *lower* GDP loss (-65.2 % vs. -90.9 %). More bankruptcies but less aggregate output damage. The resolution lies in the composition of surviving firms. In the theoretical-optimal variant, the Tobin's q investment rule drives a credit surge (credit/GDP = 2.148 vs. 0.053 in the all-best case), creating a large mass of over-levered, low-productivity firms that contribute to measured GDP while remaining financially fragile. The high ZLB frequency (99.0 %) then locks the economy into a deflationary trap, compressing GDP growth to 0.43 %. In the all-best economy, the firms that survive the bankruptcy wave have higher productivity (from the tenfold R&D increase, $\phi : 0.005 \to 0.05$) and lower leverage (from the investment cut, $\chi : 0.03$), so surviving-firm output per unit is higher even though fewer firms survive. The GDP gap is therefore smaller despite the larger bankruptcy count.

7. Discussion: Historical Case Studies

The demand-deficiency trap identified in the all-best economy is not a model artefact. Three historical episodes exhibit the same structural pattern: individually-rational conservative behaviour by households, banks, and firms simultaneously producing aggregate demand shortfalls that proved resistant to conventional monetary policy.

7.1. Japan's Lost Decade(s): 1991–2003

Japan's asset-price collapse in 1990–91 produced the clearest historical instantiation of the all-best mechanism on the household and firm dimensions. Household saving rates rose from approximately 14 % in 1990 to 17 % in 1998 as precautionary motives strengthened in response to rising unemployment and asset-price deflation [22] — the same direction as the model's $\sigma : 0.05 \to 0.20$, albeit of smaller magnitude. Japanese firms cut gross fixed capital formation by roughly 30 % between 1991 and 2002, allocating a larger share of remaining investment to product innovation (R&D) rather than capacity expansion [2] — the same investment-to-R&D substitution that the grid search identifies as individually optimal ($\chi : 0.10 \to 0.03$; $\phi : 0.005 \to 0.05$).

The bank channel in Japan diverges from the model in an instructive way. Rather than uniformly contracting credit, Japanese banks engaged in "zombie lending" — extending credit to insolvent firms to avoid recognising losses [2, 25]. This is structurally different from the model's uniform lending contraction (lending appetite: $0.80 \to 0.50$): Japanese banks were too loose toward existing insolvent clients while rationing new entrants. The aggregate demand effect was nonetheless similar — productive investment was suppressed and resources were misallocated toward non-viable firms — but the mechanism was credit misallocation rather than credit contraction. This divergence highlights a dimension of financial fragility that the model, with its representative bank, cannot capture.

The aggregate outcome mirrored the all-best prediction qualitatively: two decades of near-zero growth, rising unemployment, and persistent ZLB binding after the Bank of Japan reached the zero bound in 1999 [14]. The model's

GDP CAGR falls from 4.60 % to 1.60 %; Japan's real GDP growth averaged approximately 0.8 % per year from 1991 to 2003, against a pre-crisis trend of 4 %. Firm insolvencies rose substantially through the 1990s, peaking near three times the pre-bubble rate — smaller in magnitude than the model's 17.5-fold increase, but the model does not include the zombie-lending subsidies that Caballero et al. [2] show delayed Japanese firm exit.

One important divergence: in Japan the ZLB was associated with deflation reinforcing the ZLB through the Fisher real-interest-rate channel, while in the model ZLB frequency declines (from 80.3 % to 53.0 %) because the demand collapse eliminates inflationary pressure without generating deflation. This reflects the model's absence of asset prices and nominal debt contracts. The Japan episode therefore provides partial corroboration: the demand-deficiency mechanism operates as the model predicts, but the monetary transmission differs because of institutional features the model abstracts from.

7.2. The Euro-Area Periphery: 2010–2015

The sovereign debt crisis in Greece, Spain, and Portugal produced a second instantiation. Households in the periphery raised saving rates sharply after 2010 as unemployment rose and income uncertainty increased [26]. Spanish unemployment reached 26 % in 2013; Greek unemployment exceeded 27 %. Banks across the periphery raised capital buffers and reduced lending to the private sector under regulatory and market pressure — a pattern consistent with the historical norm that bank credit contracts sharply in the aftermath of financial crises [26]. Gross fixed capital formation fell 40–50 % in Greece and Spain between 2008 and 2013, broadly in line with the historical average capital expenditure contraction of 25–35 % documented across severe financial crises [26].

The aggregate outcome is again consistent with the model mechanism. Greece lost approximately 25 % of real GDP between 2009 and 2015; Spain lost approximately 10 %. The simultaneous austerity programmes imposed on all periphery economies by the Troika created a supra-national version of the coordination failure: each economy's individually-rational fiscal consolidation reduced aggregate demand for every other economy, and the ECB could not offset the demand shortfall at the ZLB [13]. This institutional dimension — the role of the policy framework in simultaneously incentivising caution across agent types — suggests that the model's result generalises beyond individual economies to coordination failures at the regional level.

The Euro-area episode also illustrates the distributional paradox documented in Remark 1: inequality within each country fell as high-income households cut consumption proportionally more, and household wealth positions of surviving employed workers improved relative to the crisis trough — yet aggregate welfare deteriorated sharply as the productive base eroded.

7.3. The United States Balance-Sheet Recession: 2008–2012

Koo [21] introduces the concept of a "balance-sheet recession": after an asset-price collapse, households and firms simultaneously pay down debt rather than borrow to spend, even at zero interest rates. The resulting collapse in credit demand renders conventional monetary policy ineffective, because lowering the interest rate does not stimulate new borrowing. This maps precisely onto the model's result: in the all-best economy, credit/GDP falls from 0.094 to 0.053 (a 43.4 % contraction) and ZLB frequency falls from 80.3 % to 53.0 %, not because monetary policy regained traction but because the demand for credit collapsed.

The empirical facts are consistent with this interpretation. The US household saving rate rose from approximately 2.5 % in 2005 to 6.5 % in 2009, a four-percentage-point increase driven by precautionary motives as unemployment rose to 10 % [5] — directionally identical to the model's $\sigma : 0.05 \to 0.20$. The Federal Reserve's Senior Loan Officer Survey recorded net tightening percentages above 60 % in 2008–09, analogous to the bank lending appetite falling from 0.80 to 0.50. US non-residential fixed investment fell approximately 24 % between 2007 Q4 and 2009 Q2.

Eggertsson and Krugman [13] provide the closest DSGE parallel. Their model shows that when indebted households are forced to deleverage, aggregate demand falls and the economy enters a liquidity trap. The present paper's mechanism differs in one important respect: the saving increase is *endogenous*, arising from individual optimisation rather than an exogenous debt constraint. The model therefore isolates the pure coordination failure from the balance-sheet repair mechanism: even without an initial debt shock, survival-maximising behaviour can produce the same demand-deficiency trap that Koo and Eggertsson–Krugman identify in indebted economies.

This distinction has a policy implication that the case studies collectively reinforce. In all three episodes, *conventional* (interest-rate-based) monetary policy was unable to offset the demand shortfall at the ZLB. Unconventional policies (quantitative easing, forward guidance) provided partial offset in the US and later in the Eurozone, but the transmission operated through asset prices and expectations rather than the credit channel — consistent with the model's finding that ZLB frequency declines as credit demand collapses rather than as monetary traction improves. The model's finding that coordination failure arises from survivability optimisation, not from any financial imbalance, suggests that demand-side policy interventions (fiscal stimulus, direct consumption support) are needed to shift the economy from the low-activity all-best equilibrium to the high-activity baseline, consistent with the Keynesian policy conclusion that Cooper and John [6] derive analytically from their coordination-failure framework.

8. Discussion and Conclusion

The three historical episodes examined in Section 7 confirm that the demand-deficiency mechanism identified computationally is not a model artefact. In Japan (1991–2003), the Euro-area periphery (2010–2015), and the United States (2008–2012), individually-rational conservative behaviour by households, banks, and firms simultaneously produced aggregate demand shortfalls that conventional monetary policy could not reverse at the zero lower bound. The computational model isolates the pure coordination-failure mechanism from confounding factors (nominal debt, asset prices, sovereign spreads) present in each historical episode, confirming that the trap is structural rather than episode-specific.

8.1. Two Computational Equilibria

The simulation identifies two distinct attractor states that the economy occupies under different sets of committed decision rules. The bounded-rational baseline converges to a *coordinated equilibrium* characterised by a GDP CAGR of 4.60 %, unemployment of 0.24 %, a firm bankruptcy rate of 0.27 % per year, and mean firm profit of \$4.20 per period (Table 4). The all-best economy converges to a *decentralised equilibrium* with GDP CAGR of 1.60 %, unemployment of 1.29 %, a firm bankruptcy rate of 4.73 % per year, and mean firm profit of −\$9.21. Both states are stable in the sense that neither exhibits divergent dynamics across the 50-seed Monte Carlo evaluation: the economy does not collapse to zero output or explode to unbounded growth under either rule set. The coordination failure is therefore *sustainable* — the decentralised equilibrium, while markedly inferior in aggregate welfare, is a genuine long-run attractor state, not a transitional trajectory.

Result 3. *The computational economy possesses two identifiable attractor states under committed decision rules. The coordinated equilibrium (bounded-rational baseline) is characterised by the moment vector*

$$(g^*, u^*, b^*, \pi^*) = (4.60\,\%,\ 0.24\,\%,\ 0.27\,\%/yr,\ +\$4.20).$$

The decentralised equilibrium (all-best) is characterised by

$$(g^d, u^d, b^d, \pi^d) = (1.60\,\%,\ 1.29\,\%,\ 4.73\,\%/yr,\ -\$9.21).$$

Both attractor states are stable across Monte Carlo seeds.

It is important to distinguish this characterisation from a formal Nash equilibrium. As noted in Section **??**, the all-best strategies were each optimised against the bounded-rational baseline, not against each other's best-response strategies. The decentralised equilibrium is therefore best understood as a *stochastic steady state under committed rules*: the long-run mean of the economy's aggregate statistics when agent types are locked into their survival-optimal strategies simultaneously. Whether an iterative best-response dynamic — each agent type revising its strategy against the others' updated rules — would converge to a stable Nash equilibrium, and whether that Nash equilibrium would sit closer to the decentralised equilibrium documented here or to some intermediate attractor, remains an open question addressed briefly in the future-work discussion below.

The two attractor states correspond directly to Diamond's [1982] multiplicity of equilibria in a search economy: a low-activity state in which each agent's conservative strategy is individually optimal given that all other agents are also conservative, and a high-activity state in which coordinated expansion supports greater individual welfare. The present model adds two features absent from Diamond's framework. First, the attractor states are *quantitatively characterised*: the decentralised equilibrium features GDP growth 65.2 % below the coordinated state, firm bankruptcies 17.5 times higher, and mean firm profit turning negative. Second, the transition from the coordinated to the decentralised equilibrium requires no external shock or policy shift. It results solely from each agent type independently adopting the strategy that best serves its individual survival — each of which was individually validated against the coordinated (high-activity) baseline. The selection of the low-activity attractor is therefore endogenous: it is the emergent aggregate consequence of decentralised individual optimisation, without any coordination device present to sustain the high-activity state.

Three conclusions emerge from the analysis.

Coordination failure is structural.. The all-best economy produces markedly worse *aggregate* outcomes than the bounded-rational baseline (−65.2 % GDP CAGR, +4.46 percentage points of annual firm bankruptcy, mean firm profit turning negative at −$9.21) even though every agent type is following the strategy that best serves its own survival. It bears emphasis that individual welfare is not uniformly harmed: median household wealth rises 76 % (from

\$3,842 to \$6,757), confirming that the failure is not one of individual welfare but of *social* welfare aggregated across all agent types. Households rationally defect from the coordinated consumption path; the cost is borne collectively by firms and, in the long run, by households themselves as the productive base that supports their income erodes. The failure therefore does not depend on the use of theoretically inappropriate decision rules. It is structural: a property of decentralised individual optimisation in a Keynesian economy with strategic complementarities in consumption and investment.

The failure mechanism differs from theoretical-optimal coordination failure.. In the companion paper [23], coordination failure is driven by a credit surge (Tobin's q investment) and ZLB lock-in (Euler consumption smoothing). Here, the failure operates through demand deficiency: conservative household savings and bank credit rationing suppress the consumption and investment flows that sustain firm viability. Firm bankruptcies rise 17-fold despite firms also following conservative strategies, because conservative behaviour is an individually rational response to an environment that every agent is simultaneously making more hostile. Bounded rationality therefore acts as an implicit stabiliser not merely by avoiding theoretical-optimal extremes, but by maintaining consumption and credit flows that the economy requires for internal consistency.

Implications for macro-ABM design.. The result establishes that the bounded-rational behaviour documented in heterogeneous-agent macro models [12, 15] is not a modelling convenience but a structural stabiliser. Replacing bounded rationality with survival-maximising rules — even without imposing theoretical equilibrium conditions — is sufficient to produce coordination failure. Future work should investigate whether an iterative best-response dynamic (each agent type updating its strategy in response to others' strategies) converges to a stable Nash equilibrium or cycles, and whether the resulting equilibrium is closer to the bounded-rational baseline or to the all-best outcome documented here.

Acknowledgements

Simulation code available at the project repository.

Conflict of interest

The author declares no conflict of interest.

References

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

[2] Caballero, R.J., Hoshi, T., Kashyap, A.K., 2008. Zombie lending and depressed restructuring in Japan. American Economic Review 98, 1943–1977.

[3] 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.

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

[5] Carroll, C.D., Slacalek, J., Tokuoka, K., White, M.N., 2017. The distribution of wealth and the marginal propensity to consume. Quantitative Economics 8, 977–1020.

[6] Cooper, R., John, A., 1988. Coordinating coordination failures in Keynesian models. Quarterly Journal of Economics 103, 441–463.

[7] 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.

[8] Delli Gatti, D., Gallegati, M., Greenwald, B., Russo, A., Stiglitz, J.E., 2010. The financial accelerator in an evolving credit network. Journal of Economic Dynamics and Control 34, 1627–1650.

[9] Diamond, P.A., 1982. Aggregate demand management in search equilibrium. Journal of Political Economy 90, 881–894.

[10] Dosi, G., Fagiolo, G., Napoletano, M., Roventini, A., 2013. Income dis-

tribution, credit and fiscal policies in an agent-based Keynesian model. Journal of Economic Dynamics and Control 37, 1598–1625.

[11] 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.

[12] 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.

[13] Eggertsson, G.B., Krugman, P.R., 2012. Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach. Quarterly Journal of Economics 127, 1469–1513.

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

[15] 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.

[16] Fisher, I., 1933. The debt-deflation theory of great depressions. Econometrica 1, 337–357.

[17] Guerrieri, V., Lorenzoni, G., 2017. Credit crises, precautionary savings, and the liquidity trap. Quarterly Journal of Economics 132, 1427–1467.

[18] Heller, W.P., 1986. Coordination failure under complete markets with applications to effective demand, in: Heller, W.P., Starr, R.M., Starrett, D.A. (Eds.), Equilibrium Analysis: Essays in Honor of Kenneth J. Arrow. Cambridge University Press, Cambridge. volume 2, pp. 155–175.

[19] Keynes, J.M., 1936. The General Theory of Employment, Interest and Money. Macmillan, London.

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

[21] Koo, R.C., 2011. The world in balance sheet recession: Causes, cure, and politics. Real-World Economics Review , 19–37.

[22] 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.

[23] Leovonzko, E., 2026. Theoretical optimality and collective failure in a Keynesian agent-based model. Journal of Economic Dynamics and Control Working paper.

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

[25] Peek, J., Rosengren, E.S., 2005. Unnatural selection: Perverse incentives and the misallocation of credit in Japan. American Economic Review 95, 1144–1166.

[26] Reinhart, C.M., Rogoff, K.S., 2009. This Time Is Different: Eight Centuries of Financial Folly. Princeton University Press, Princeton, NJ.

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