

Structural Dualism, Under-Development Traps in a Developed Economy: The case of the Irish Economy

Srinivas Raghavendra

J.E. Cairnes School of Business and Economics, University of Galway

s.raghav@universityofgalway.ie

PAPER SUBMITTED TO THE FMM CONFERENCE 2026

Abstract

Ireland's national accounts describe two economies in one jurisdiction: a foreign-owned multinational enclave with labour productivity six to eight times that of the domestic sector, and an indigenous sector that supplies most employment but little measured output growth. Can this structure trap the domestic economy at low growth? We show that the classical Thirlwall mechanism cannot be the answer for a current-account-surplus member of a currency union, and build a *linkage-constrained* growth model in which the enclave plays the role of the external sector, domestic productivity follows Verdoorn's law, and domestic wages are pulled towards enclave wage levels. In the linear model the only instability is monotonic divergence at a Harrodian knife-edge; adding a Kaleckian consumption channel removes it in a wage-led economy but leaves it in a profit-led one. A genuine trap, with a stable low state, a stable high state and an unstable tipping point between them, requires nonlinearity, which we supply with a kinked investment accelerator, a ceiling on domestic growth, and the wage share as a slow third state. The model generates hysteresis: each year in the trap raises the wage share and the investment hurdle, so the same policy push works early and fails late. We show which fiscal levers end the trap, why internal devaluation changes the path but not the destination, and which assumptions the result does and does not depend on, including whether dualism itself is assumed.

JEL Classification: O11, F43, E12, E32, E62

Keywords: structural dualism, balance-of-payments-constrained growth, cumulative causation, multiple equilibria, hysteresis, Ireland

1. Introduction

Ireland's national accounts describe two structurally distinct economies. One is a small number of foreign-owned multinational enterprises (MNEs), concentrated in pharmaceuticals, medical devices, and information and communication technology, whose value added per hour worked is roughly six to eight times that of the rest of the economy. The other is the indigenous economy in which most residents work, earn and spend. This is not a recent statistical curiosity. The divergence has widened since the mid-2000s and was thrown into sharp relief by the 2015 national-accounts revision, which raised measured GDP growth for that single year to 26 per cent. The Central Statistics Office's foreign/domestic productivity split for the second quarter of 2025 puts labour productivity at €530.90 per hour in MNE-dominated sectors against €63.40 per hour elsewhere, and shows the two series moving in opposite directions from quarter to quarter (Central Statistics Office, 2025). The OECD reports multinational-sector productivity growth of 10.9 per cent against 2.5 per cent for indigenous firms in 2017 (OECD, 2026). As a result Ireland records the highest GDP per hour worked in the OECD, above the United States, Norway and Denmark, a fact that reflects the location of intellectual property and contract manufacturing rather than the productive capacity of resident workers (OECD, 2026).

We call this structure *structural dualism*: an enclave whose growth is exogenous to the domestic economy and whose linkages to it are weak. It is not dualism in the Lewisian sense of a labour surplus awaiting absorption. This paper asks whether such a structure can trap the *domestic* economy at low growth, and if so through which mechanism.

The natural starting point is Thirlwall's balance-of-payments-constrained growth model (Thirlwall, 1979, 2011), in which growth is limited by export growth divided by the income elasticity of imports. Inside the euro area, however, Ireland no longer faces a foreign-exchange constraint of the usual kind. TARGET2 settles cross-border euro payments through central bank balance sheets rather than requiring a country first to acquire the currency needed to pay for its imports. Persistent imbalances are therefore no longer disciplined by the exhaustible financing (reserves, foreign borrowing, a self-correcting depreciation) that makes Thirlwall's constraint bind (Sinn and Wollmershäuser, 2012; Cesaratto, 2013). The 2010–2012 euro-area crisis put this to the test. Private capital flows to peripheral economies, Ireland among them, suddenly stopped, which would force an adjustment under Thirlwall's model, but Eurosystem financing filled the gap.

The external position since then points the same way. The current account is in persistent surplus (between 8 and 16 per cent of GDP a year over 2021–25; €24.3 billion, or 4.0 per cent of GDP, in 2025 on the modified CA* measure; Central Statistics Office, 2026a,b), and the dominant export sector responds to global corporate decisions rather than to Irish prices or demand. Recent policy analysis agrees. The IMF projects the current account surplus to stay at 7 to 8 per cent

of GDP through 2029 and judges the external position moderately stronger than fundamentals imply. It also describes corporation tax receipts as highly concentrated, and recommends that excess receipts go to the two national savings funds and that the medium-term spending plan act as a binding ceiling (International Monetary Fund, 2026). The Fiscal Council reports that only one euro in six of corporation tax receipts is set aside, that excluding excess corporation tax the budget would be in deficit by €11 billion (3 per cent of GNI*) in 2026, and that planned net spending growth is the fastest in the European Union (Irish Fiscal Advisory Council, 2026). On both accounts the binding question for Ireland is not how to finance imports but how the enclave's receipts and demand reach, and are used by, the domestic economy.

Yet the absence of a foreign-exchange constraint does not mean the absence of any external demand constraint. Harrod's logic still requires a source of demand growth that the domestic economy depends on but does not itself generate. For Ireland that role is filled not by the rest of the world in the conventional balance-of-payments sense but by the foreign-owned enclave domiciled in Ireland. The enclave is external in every structural sense that matters, though domestic in the national accounts: an open-economy constraint operating inside what the accounts register as a single economy. The recursion is almost fractal. The Harrod–Thirlwall logic that no longer binds between Ireland and the rest of the world reappears one level down, between the enclave and the domestic economy. The binding constraint becomes the enclave's linkage with the domestic economy (wages spent locally, corporation tax recycled into public spending, local procurement), weighed against the domestic economy's own propensity to spend on imports. We call the resulting relation the *linkage-constrained growth condition*. It is formally analogous to Thirlwall's law, with enclave output growth in place of export growth and a linkage parameter in place of the import elasticity. The parameter reflects the weakness of backward linkages between the multinational and indigenous sectors, so domestic growth is limited by the share of enclave growth that actually reaches the domestic economy.

The paper builds on Thirlwall's framework, itself an extension of Harrod's foreign-trade multiplier (Harrod, 1933). We embed the linkage condition in a Kaldorian cumulative-causation structure in which domestic productivity growth follows Verdoorn's law (Verdoorn, 1949; Kaldor, 1966) and domestic wage growth is partly set by contagion from the much higher wage level in the enclave. This is consistent with recent ESRI evidence that pay pressure in the foreign sector spills into the domestic sector through shared labour and housing markets (ESRI, 2025). It is also consistent with the IMF's finding that, while Ireland shows little evidence of classical Dutch-disease resource-movement or spending effects, productivity spillovers from multinationals to indigenous firms have been limited and linkages between the two remain weak (International Monetary Fund, 2023).

The paper makes five contributions.

1. **A linkage-constrained cumulative-causation model.** The linkage condition is combined with Verdoorn’s law and a wage equation in which domestic wages catch up with enclave wage levels. This yields a two-state dynamical system in domestic growth and the log wage gap with a closed-form stability condition (Section 4).
2. **The role of distribution.** In the linear model the only instability is *monotonic divergence* at a Harrodian knife-edge; it is not a trap. Adding a Kaleckian consumption channel shows that the knife-edge exists in a profit-led economy and disappears in a wage-led one, where damped cycles replace it (Section 5).
3. **A genuine trap, with hysteresis.** A stable low state, an unstable tipping point and a stable high state require nonlinearity. We obtain them from a kinked investment accelerator, a composite ceiling on domestic growth and the wage share as a slow third state, derive the three equilibria and their stability in closed form, and show that time spent in the trap raises the barrier to escape (Section 6).
4. **Whether dualism is generated or assumed.** We ask whether the result depends on dualism being put in by hand, on domestic wages following the enclave, or on the profit-led regime, and report the answers (Sections 7, 6.10).
5. **Policy implications for Ireland.** We show which fiscal levers can end the trap and which cannot, why timing matters (the same push works early and fails late), and why an internal devaluation changes the path but not the destination (Sections 8, 9). We also compare procyclical, flat and countercyclical fiscal rules under an enclave cycle. The same downturn that a countercyclical rule absorbs costs the economy its high state under a flat or procyclical rule, which bears on the current debate over whether windfall corporation tax receipts should be saved or spent (Section 8.1).

The rest of the paper is organised as follows. Sections 4–6 build the model: the base model, the Kaleckian block and the nonlinear model with a trap. Section 7 tests how much depends on dualism being assumed. Sections 8 and 9 analyse fiscal levers, timing and rules, and apply the model to Ireland after 2008. Section 10 relates the model to Thirlwall’s law, Section 11 lists limitations and Section 12 concludes.

2. Stylised Facts

Four facts motivate the model.

Fact 1: The productivity gap is large and has recently been widening rather than closing. CSO data for 2025 Q2 show domestic-sector multifactor productivity falling by 2.4 percent quarter-

on-quarter while foreign-sector multifactor productivity rose by 2.1 percent (Central Statistics Office, 2025). This is not a single-quarter anomaly: OECD figures for 2017 show foreign-owned firm productivity growth of 10.9 percent against 2.5 percent for indigenous firms (OECD, 2026), consistent with a persistent divergence documented over the post-2015 period by ESRI (2025).

Fact 2: The two sectors are connected through the labour market, not primarily through supply chains. The IMF finds little evidence of the classical Dutch-disease channels but does find that productivity spillovers from multinationals to the indigenous sector have been limited, and recommends fostering MNE–domestic linkages, noting the absence of the deep supply-chain relationships found in, for example, the German or Japanese automotive sectors (International Monetary Fund, 2023). ESRI (2025) document that MNEs pay persistently higher wages than domestic firms and argue that this creates pay pressure in the domestic sector through shared labour-market competition and, indirectly, through housing costs, a channel that operates independently of any commercial relationship between the sectors. We note that this evidence on wage catch-up is indirect: a direct, per-employee foreign-versus-domestic wage series on the CSO’s sector classification is still to be assembled.

Fact 3: Measured aggregate output overstates the resources available to the resident economy. The gap between GDP and Modified Gross National Income (GNI*), the CSO measure designed to strip out re-domiciled company profits and the depreciation on re-domiciled intellectual property and aircraft leasing, was such that 2016 GDP was 143 percent of GNI*. A bottom-up decomposition of national income finds that domestic sectors’ income runs on average 4.5 percentage points below GNI* itself over 2013–2021, reflecting retained foreign-sector savings that never enter domestic circulation (Irish Fiscal Advisory Council, 2023).

Fact 4: The foreign sector’s scale bears little relation to its employment or linkage footprint. Foreign-owned multinationals account for a large share of corporation tax revenue and roughly a quarter of private-sector employment, but their capital- and IP-intensity decouples value-added growth from employment growth and from the volume of local procurement, limiting the mechanical size of any linkage effect on domestic demand (International Monetary Fund, 2023).

Together these facts describe an economy in which the productivity gap is wide and not visibly closing, the principal transmission channel between sectors runs through wages and property costs rather than commercial linkages, and conventional aggregates overstate the resources available to fund domestic growth. The model is built around exactly these features.

3. Related Literature

The paper draws on five literatures that have developed largely separately.

Balance-of-payments-constrained growth. Thirlwall’s law, $g = x/\pi$, holds that long-run

growth is bounded by export growth divided by the income elasticity of import demand, since faster growth would generate an unsustainable current-account deficit (Thirlwall, 1979). The model synthesises Harrod’s foreign-trade multiplier, Prebisch’s centre–periphery structuralism and the foreign-exchange constraint of Chenery’s two-gap model (Harrod, 1933; Thirlwall, 2011). Thirlwall and Hussain (1982) show that capital inflows relax the constraint only temporarily. Thirlwall (2011) also discusses a self-reinforcing underdevelopment trap in which specialisation in low-income-elasticity goods depresses growth and discourages the structural change that would relax the constraint. We take the logic of that trap to the enclave–domestic boundary, where the transmission mechanism is different.

Cumulative causation. Kaldor (1966) argued that manufacturing growth drives economy-wide productivity growth through increasing returns, formalised in Verdoorn’s law (Verdoorn, 1949). An initial growth advantage improves competitiveness and reinforces itself, in contrast to the neoclassical presumption of convergence. We use this loop for the domestic sector and add the wage-level gap to the enclave as a second feedback.

Distribution and demand regimes. Kalecki (1971) established that wage earners spend a larger share of income than profit recipients, so that distribution has independent consequences for demand. Bhaduri and Marglin (1990) formalised the resulting ambiguity: demand can respond positively or negatively to a rising wage share depending on whether the consumption gain or the investment loss dominates. We embed the distinction in a dual-economy setting in which a large share of profit income leaks abroad.

Harroddian instability, floors and ceilings. Harrod (1939) showed that an accelerator–multiplier economy sits on a knife-edge, and Hicks (1950) and Goodwin (1951) showed that a floor (investment cannot be negative) and a ceiling (output is limited by capacity) bound the instability and generate fluctuations. Our nonlinear model is in this tradition, with the ceiling interpreted as an external, supply or fiscal limit on the domestic sector. The supermultiplier literature (Serrano, 1995; Freitas and Serrano, 2015) restores stability by letting the accelerator adjust capacity to demand; we leave that adjustment out and ask what the floor and the ceiling alone generate. Goodwin (1967) is the natural next step for the wage share and employment (Section 11).

Poverty traps, and Ireland’s dual economy. The possibility of a stable depressed equilibrium coexisting with a feasible higher one has a long history after Nelson (1956); Azariadis and Stachurski (2005) survey it. Applied work on Ireland has concentrated on measurement (GNI* and related aggregates: Irish Fiscal Advisory Council, 2023), on the size and persistence of the domestic–foreign productivity gap (ESRI, 2025; Central Statistics Office, 2025), and on whether Ireland shows Dutch-disease effects from its FDI boom (International Monetary Fund, 2023). To our knowledge this literature has not been formalised as a growth model with a trap.

4. The Base Model: Linkage-Constrained Cumulative Causation

4.1 The idea in a nutshell

Start with what the model is trying to capture. Ireland's enclave grows because of decisions taken elsewhere: where to book intellectual property, where to locate production, where to pay tax. The domestic economy does not control that growth, but it receives some of it: through wages spent locally, through corporation tax recycled into public spending and through local procurement. Some of what the domestic economy earns then leaks out again, as imports and as profit repatriation. How fast the domestic economy grows depends on how much of the enclave's growth reaches it net of leakage, and on whether domestic growth can feed itself.

It can, through a familiar loop. Faster growth raises productivity (Verdoorn's law: more output allows more specialisation and learning by doing). Wages follow productivity only partly, so unit labour costs fall relative to those of the enclave, which is the benchmark here for what domestic firms can pay and still compete. Lower relative costs raise demand, and demand raises growth again. This is cumulative causation (Kaldor, 1966). But growth also changes the wage-level gap between the sectors: when domestic productivity rises, wages rise, and the gap that was pulling domestic wages up towards enclave levels narrows. Two things therefore move together, domestic growth and the wage gap, and the question for stability is whether the loop through them dies out or builds on itself.

The only thing that can go wrong in such a linear model is a *runaway*: a shock that is amplified instead of absorbed. Whether it happens depends on the strength of the accelerator-multiplier, σ , and on how strongly the cost-competitiveness loop feeds back on demand, Λ ; the speed of wage catch-up ϕ affects how long adjustment takes but not whether it succeeds. Near the boundary the economy converges so slowly that it looks stuck, long before it is formally unstable. That is the first, linear approximation to the persistence we later explain with a genuine trap.

4.2 Environment and assumptions

Time is discrete, one period is a year. Lower-case letters denote growth rates (in per cent a year) of the corresponding upper-case level. An economy has an enclave sector F and a domestic sector D .

Definition 1 (Structural dualism). *An economy is structurally dual if it contains an enclave sector whose output growth y_F is determined outside the domestic economy and whose linkages to the domestic sector (local procurement, wages spent locally, tax revenue recycled into domestic demand) are weak relative to its size. This is not a Lewisian labour-surplus dualism.*

Assumption 1. *Enclave output growth $y_{F,t}$ is determined by global corporate decisions on location, taxation and licensing, and is treated as exogenous. Enclave wage growth ω_F^0 is likewise*

exogenous from the domestic sector's point of view: it is set by enclave productivity and international pay norms.

Assumption 2. *Domestic output is demand-determined: production adjusts to effective demand for domestically produced goods and services, as in the post-Keynesian tradition underlying both Harrod's multiplier and Thirlwall's law.*

The one log *level* in the model is Γ_t , the wage-level gap between the enclave and the domestic sector. The model has six equations.

Demand. Domestic growth has three sources: momentum, the linkage to the enclave, and autonomous demand $g_{A,t}$, of which a share χ leaks into imports or profit repatriation:

$$q_{D,t} = \sigma q_{D,t-1} + \ell y_{F,t} + (1 - \chi) g_{A,t}. \quad (1)$$

Here σ is a Harrodian accelerator-multiplier coefficient: the share of last year's growth that carries over into this year's, through induced spending and investment. The term $\ell = \theta \kappa$ is the pass-through of enclave growth into domestic demand. The term ℓy_F is the linkage-constrained analogue of Thirlwall's x/π : in his model exports pull the economy and imports leak demand; here the enclave pulls and the linkage parameter measures how much of the pull reaches the domestic economy.

Verdoorn's law and wages.

$$p_{D,t} = a + b q_{D,t}, \quad (2)$$

$$w_{D,t} = \zeta p_{D,t-1} + \phi \Gamma_{t-1} + \psi h_t, \quad \zeta, \phi, \psi \in (0, 1). \quad (3)$$

Domestic productivity growth rises with output growth (slope b). Domestic wage growth has three drivers: a share ζ of lagged productivity growth, a catch-up term $\phi \Gamma_{t-1}$ that pulls domestic wages towards the enclave level, and housing and living-cost growth h_t . The catch-up term is the wage-contagion channel, a behavioural assumption: workers and employers in the domestic sector look at what enclave workers earn and share housing and labour markets with them, so they cannot let the gap grow indefinitely. Section 7 tests how far the results depend on it.

Competitiveness, feedback and the wage gap. Relative unit labour cost growth is $\text{rulc}_{D,t} \equiv w_{D,t} - p_{D,t-1} = \phi \Gamma_{t-1} - (1 - \zeta) p_{D,t-1} + \psi h_t$. It feeds back on autonomous demand, and the gap

evolves with the difference between enclave and domestic wage growth:

$$g_{A,t} = g_A^0 - \eta \text{rulc}_{D,t}, \quad \eta > 0, \quad (4)$$

$$\Gamma_t = \Gamma_{t-1} + \omega_F^0 - w_{D,t}. \quad (5)$$

Faster growth raises productivity (Verdoorn), lowers rulc_D through $-(1 - \zeta)p_D$, raises autonomous demand and so raises growth again. Note that a falling rulc_D involves no wage cut: when productivity rises by p_D , wages rise by ζp_D , which is less than p_D because $\zeta < 1$. In this base version higher wages matter only as a cost: they erode competitiveness and so reduce demand, as in the standard competitiveness view of wages. Section 5 adds their role as income.

4.3 The two-equation system

Only q_D and Γ have lagged values on the right-hand sides, so they are the two states. Define the loop gain $\Lambda \equiv (1 - \chi)\eta b$: the amount by which a one-point rise in growth raises next year's growth by way of productivity, relative costs and demand, net of leakage. Substituting (2)–(5) into (1) (Appendix A) gives

$$\begin{pmatrix} q_{D,t} \\ \Gamma_t \end{pmatrix} = \begin{pmatrix} \mu_{1,t} \\ \mu_{2,t} \end{pmatrix} + \begin{pmatrix} \tilde{\sigma} & -\beta \\ -\gamma & \delta \end{pmatrix} \begin{pmatrix} q_{D,t-1} \\ \Gamma_{t-1} \end{pmatrix}, \quad (6)$$

with $\tilde{\sigma} = \sigma + (1 - \zeta)\Lambda$, $\beta = (1 - \chi)\eta\phi$, $\gamma = \zeta b$, $\delta = 1 - \phi$, and forcing terms

$$\mu_{1,t} = \ell y_{F,t} + (1 - \chi)[g_A^0 + \eta(1 - \zeta)a - \eta\psi h_t], \quad \mu_{2,t} = \omega_F^0 - \zeta a - \psi h_t.$$

Table 1 explains the economic meaning of each matrix entry.

Entry	Expression	What it says
$\tilde{\sigma}$	$\sigma + (1 - \zeta)\Lambda$	<i>Growth feeds on itself.</i> Last year's growth carries over through the accelerator (σ) and through the part of the cost-competitiveness loop that works directly on growth. Larger than σ : cumulative causation adds to the accelerator.
$-\beta$	$-(1 - \chi)\eta\phi$	<i>A wide wage gap holds growth down.</i> A large wage gap means strong catch-up pressure on domestic wages, hence higher relative costs and weaker competitiveness.
$-\gamma$	$-\zeta b$	<i>Growth narrows the wage gap.</i> Faster growth raises productivity, hence wages, so the domestic sector catches up.
δ	$1 - \phi$	<i>The wage gap is persistent.</i> Only a fraction ϕ of the gap is closed each year by catch-up.

Table 1: Economic meaning of the matrix entries.

The forcing terms μ_t contain the time series y_F and h_t , but *stability depends only on the matrix*. With $y_F = y_F^*$ and h constant there is a steady state,

$$q_D^* = \frac{\ell y_F^* + (1 - \chi)[g_A^0 - \eta(\omega_F^0 - a)]}{1 - \sigma - \Lambda}, \quad (7)$$

which does not depend on ϕ , ζ or ψ : what matters is enclave wage growth relative to the autonomous productivity trend, $\omega_F^0 - a$. The intuition is that in a steady state the gap is constant, so domestic wages must grow at the enclave rate ω_F^0 . If the domestic productivity trend a falls short of that rate, domestic unit labour costs rise relative to the enclave's by $\omega_F^0 - a$ a year, and the competitiveness effect $\eta(\omega_F^0 - a)$ takes that much away from growth. The wage gap itself, Γ^* , scales as $1/\phi$: the weaker the catch-up pull, the larger the gap needed to deliver the same wage growth.

4.4 Stability and the Harroddian knife-edge

Nullclines and the phase diagram. Write $\hat{q} = q_D - q_D^*$ and $\hat{\Gamma} = \Gamma - \Gamma^*$ for deviations from the steady state. A *nullcline* is the set of points at which one variable stops changing. The growth nullcline $\Delta q_D = 0$ is $\hat{\Gamma} = -\frac{1-\tilde{\sigma}}{\beta}\hat{q}$; the gap nullcline $\Delta \Gamma = 0$ is $\hat{\Gamma} = -\frac{\gamma}{\phi}\hat{q}$. Both slope downward. Along the first, a wider gap just offsets higher growth, so that next year's growth equals this year's; along the second, higher growth just offsets a wider gap, so that the gap stays put. The two downward slopes are the loop in geometric form: growth and the gap pull in opposite directions on each other, and in the same direction on themselves.

Figure 1 shows three economies that differ only in the accelerator σ .

Phase diagrams of the base model in deviations from the steady state (open circle = start of a trajectory)




Figure 1: Phase diagrams of the base model in deviations from the steady state, for $\sigma = 0.60, 0.88$ and 0.92 ($b = 0.5$, $\chi = \eta = 0.3$, $\zeta = 0.85$, $\phi = 0.2$, so that $\Lambda = 0.105$ and the knife-edge is at $\sigma = 0.895$). Orange: $\Delta q_D = 0$. Blue: $\Delta \Gamma = 0$. Black: eigen-directions, with arrowheads pointing inward where deviations shrink and outward where they grow. Thin lines with dots: trajectories from four starting points (open circles). At $\sigma = 0.60$ the nullclines cross at a wide angle and trajectories converge quickly. At $\sigma = 0.88$ they are almost parallel and trajectories first move to them and then creep along. At $\sigma = 0.92$ the slow eigenvalue exceeds one: a saddle.

Economic interpretation. In the left panel the economy is comfortably stable. A shock to growth is absorbed quickly: the nullclines cross at a wide angle, trajectories head straight for them and then slide in to the steady state. In the middle panel the accelerator is stronger. The nullclines are nearly parallel, so after the quick initial adjustment the economy moves *along* them, very slowly, towards the steady state. Economically, a growth shock leaves behind a wage gap, and the gap itself holds growth down, which in turn keeps the gap open: the loop makes the shock persist. In the right panel the slow eigenvalue exceeds one and the steady state is a saddle: the economy can converge only along a single stable path, and off it, a growth shortfall deepens year after year instead of being absorbed. No trajectory spirals. The eigenvalues are real, because both propagation channels have the same sign (growth lowers the gap, the gap lowers growth, and the product of the two off-diagonal entries is positive), so deviations move along the two eigen-directions and never rotate round the steady state.

The algebra of stability. The Schur–Cohn conditions for a 2×2 system are $p(1) = 1 - T + D > 0$, $p(-1) = 1 + T + D > 0$ and $|D| < 1$, where T and D are the trace and determinant. Using

$\beta\gamma = \Lambda\zeta\phi$ (Appendix A),

$$p(1) = \phi[1 - \sigma - \Lambda], \quad p(-1) = (1 + \tilde{\sigma})(2 - \phi) - \Lambda\zeta\phi, \quad D = \tilde{\sigma}\delta - \Lambda\zeta\phi. \quad (8)$$

The discriminant is $T^2 - 4D = (\tilde{\sigma} - \delta)^2 + 4\Lambda\zeta\phi \geq 0$, so the eigenvalues are real and there are no cycles. The nullclines are parallel exactly when $(1 - \tilde{\sigma})\phi = \beta\gamma$, that is $p(1) = 0$, which is why the knife-edge shows up in the figure as parallel lines.

Proposition 1 (Base model). *The base system is stable if and only if $\sigma + \Lambda < 1$, i.e. $\sigma < 1 - (1 - \chi)\eta b$. The conditions $p(-1) > 0$ and $|D| < 1$ cannot fail before $p(1) > 0$ does. Failure is monotonic divergence: a saddle with one eigenvalue above one.*

Harrodian knife-edge: generated, not assumed. In Harrod's model the warranted path is a knife-edge because induced investment pushes an economy that is above (below) the warranted rate further above (below) it. Taken alone, the accelerator here is stable if and only if $\sigma < 1$, and we maintain $\sigma < 1$. The instability that remains has a different source. The Verdoorn–competitiveness loop is a second, *reinforcing* mechanism that adds to the accelerator: the effective accelerator is $\sigma_{\text{eff}} = \sigma + \Lambda$, and the economy is stable if and only if $\sigma_{\text{eff}} < 1$. Cumulative causation can thus manufacture a Harrod-type instability from an accelerator that would be stable on its own. Harrod's knife-edge is an assumption about σ ; here it is a result of the dual economy's feedback structure. The threshold involves neither the wage pass-through ζ nor the catch-up speed ϕ : the first only splits the loop between the two equations, and the second sets only the speed of adjustment. A stability condition free of the catch-up speed is a statement about *structure* (accelerator, leakage, competitiveness elasticity, Verdoorn slope), not about wage-setting behaviour.

The multiplier and dynamics near the knife-edge. The steady-state multiplier $1/(1 - \sigma_{\text{eff}})$ has the form of the Keynesian multiplier (Kahn, 1931; Keynes, 1936), with σ_{eff} in the role of the marginal propensity to spend. As the boundary is approached it explodes, and so does the time the economy takes to settle. The values below are computed at illustrative parameter values (those of Figure 1, which give $\Lambda = 0.105$), not calibrated to Irish data:

σ	Multiplier $1/(1 - \sigma - \Lambda)$	Dominant eigenvalue	Half-life (years)
0.30	1.7	0.834	3.8
0.60	3.4	0.870	5.0
0.80	10.5	0.942	11.5
0.88	66.7	0.990	68
0.894	1000	0.9993	≈ 1000
0.90	-200	1.0035	unstable (saddle)

Parameters as in Figure 1 ($\Lambda = 0.105$). Past the boundary the multiplier changes sign and the steady-state formula is meaningless.

The table shows three zones. Well inside the boundary the economy is stable and adjusts quickly. Close to boundary the economy is formally stable but outcomes become hypersensitive to small changes in parameters or shocks, and a one-off shock can have effects lasting decades: a first, linear hint of an economy that looks stuck. Beyond it there is a saddle. The linear model has one equilibrium, so it cannot produce a trap: its instability is *monotonic divergence*, a different phenomenon. A genuine trap needs nonlinearity (Section 6).

What the base model leaves out. The only feedback from growth to demand in this model is reinforcing. Faster growth lowers relative unit labour costs and the wage share, and in this version a lower wage share always raises demand. Nothing allows faster growth to erode its own demand, which is the equilibrating mechanism missing from Harrod’s model. Post-Keynesian theory supplies it in different ways: Kaldor through distributive adjustment, Kaleckian models through capacity utilisation. With a single distributive channel that works only in the destabilising direction, the cost-only base model amplifies disturbances instead of damping them. The next section adds the Kaleckian channel, and Section 6 adds an investment function in a nonlinear setting.

5. Adding Distribution: The Kaleckian Block

5.1 The missing channel: wages are costs, and also incomes

In the base model a rise in rulc_D can only hurt demand. But rulc_D is also the growth of the domestic wage share, and wages are household income as well as firms’ costs. When the wage share rises, firms in the domestic sector are less competitive, but wage-earners have more to spend, and they

spend a larger part of it locally than profit-earners, a share of whose income is repatriated. Whether the net effect on demand is negative or positive is exactly the distinction between *profit-led* and *wage-led* demand regimes in the Kaleckian tradition (Kalecki, 1971; Bhaduri and Marglin, 1990).

Profit-led and wage-led regimes. Whether faster productivity growth feeds back positively or negatively on demand depends on how demand responds to the wage share. Suppose productivity rises faster than wages, so the wage share falls. In a profit-led regime, the gain in competitiveness and profitability outweighs the loss of household income: demand and growth rise, productivity rises further through Verdoorn effects, and the wage share falls again. The feedback is reinforcing.

profit-led: $\downarrow \text{wage share} \Rightarrow \uparrow \text{demand} \Rightarrow \uparrow \text{growth} \Rightarrow \uparrow \text{productivity} \Rightarrow \downarrow \text{wage share}$ (reinforcing)

In a wage-led regime, the fall in household spending outweighs the competitiveness gain: demand and growth fall, productivity growth slows, and the wage share recovers. The feedback is stabilising. The sign of this single response therefore decides whether the economy diverges or converges, and it is the equilibrating channel that Harrod's model and our base model lack.

wage-led: $\downarrow \text{wage share} \Rightarrow \downarrow \text{demand} \Rightarrow \downarrow \text{growth} \Rightarrow \downarrow \text{productivity} \Rightarrow \uparrow \text{wage share}$ (stabilising)

5.2 Two channels, one variable

Let a share s_C of autonomous demand be exposed to the consumption channel and the rest to the competitiveness channel. Households consume c_w of each unit of wage income and c_π of each unit of profit income, and a share ρ of profit is repatriated, so the excess propensity to spend out of wages is $\xi_0 \equiv c_w - c_\pi(1 - \rho) > 0$. A shift of income from profits to wages then raises consumption demand by ξ_0 per unit shifted. Replacing (4) by the weighted average of the two channels gives

$$g_{A,t} = g_A^0 - \eta_{\text{net}} \text{rulc}_{D,t}, \quad \eta_{\text{net}} \equiv (1 - s_C) \eta - s_C \xi_0. \quad (9)$$

The extension replaces η by η_{net} and changes nothing else, but η_{net} can take either sign. When $\eta_{\text{net}} > 0$ the competitiveness effect dominates and the economy is *profit-led*. When $\eta_{\text{net}} < 0$ the consumption effect dominates and the economy is *wage-led*. The boundary is $s_C^* = \eta / (\eta + \xi_0)$. Because ρ is large in an MNE-hosting economy, ξ_0 is plausibly larger than in a closed economy, which pushes towards the wage-led side. A competitiveness channel that runs mainly through the enclave weakens η for the domestic sector and points the same way. This is a conjecture to be tested, not a result.

With $\Lambda_{\text{net}} = (1 - \chi) \eta_{\text{net}} b$, all expressions of Section 4 hold with Λ_{net} in place of Λ .

5.3 What changes in the matrix, and why everything follows from one sign

The Kaleckian extension leaves the structure of the two-state system unchanged. It replaces η by η_{net} in the first row, so the system keeps the form of (6):

$$\begin{pmatrix} q_{D,t} \\ \Gamma_t \end{pmatrix} = \begin{pmatrix} \mu_{1,t} \\ \mu_{2,t} \end{pmatrix} + \begin{pmatrix} \tilde{\sigma}_{\text{net}} & -\beta_{\text{net}} \\ -\gamma & \delta \end{pmatrix} \begin{pmatrix} q_{D,t-1} \\ \Gamma_{t-1} \end{pmatrix}, \quad (10)$$

where

$$\tilde{\sigma}_{\text{net}} = \sigma + (1 - \zeta)\Lambda_{\text{net}}, \quad \beta_{\text{net}} = (1 - \chi)\eta_{\text{net}}\phi, \quad \Lambda_{\text{net}} = (1 - \chi)\eta_{\text{net}}b, \quad (11)$$

and $\gamma = \zeta b$ and $\delta = 1 - \phi$ are unchanged. Because η_{net} can be negative, the signs of Λ_{net} and β_{net} can reverse, and with them the stability properties of the system. The table shows how each entry changes.

Only the first row of the matrix changes, through $\eta \rightarrow \eta_{\text{net}}$:

Entry	Expression	Profit-led ($\eta_{\text{net}} > 0$)	Wage-led ($\eta_{\text{net}} < 0$)
$\tilde{\sigma}$	$\sigma + (1 - \zeta)\Lambda_{\text{net}}$	$> \sigma$	$< \sigma$
$-\beta_{\text{net}}$	$-(1 - \chi)\eta_{\text{net}}\phi$	negative	positive
$-\gamma$	$-\zeta b$	negative	negative
δ	$1 - \phi$	positive	positive

In the profit-led regime the two off-diagonal entries have the same sign, as in the base model: growth and the gap pull on each other in opposite directions. In the wage-led regime they have *opposite* signs: a wide gap now *raises* growth (it signals high wages, hence high consumption), while faster growth still narrows the gap. That sign change is the whole difference between the two regimes. It turns a chain of mutually reinforcing movements into one in which each variable pushes the other round a circle.

5.4 The phase diagrams: the sign of η_{net} decides the fate

Figure 2 repeats the exercise of Figure 1, now at a single value of the accelerator, $\sigma = 0.90$, and varies only the sign and size of η_{net} .

Phase diagrams at $\sigma = 0.90$: only the sign of η_{net} changes (open circle = start; arrows give the direction of motion)




Figure 2: Phase diagrams at $\sigma = 0.90$ for three economies ($b = 0.5$, $\chi = \eta = 0.3$, $\zeta = 0.85$, $\phi = 0.2$; the Kaleckian cases use $(s_C, \xi_0) = (0.2, 0.3)$ and $(0.7, 0.8)$). Left: the cost-only base model is past its knife-edge ($\sigma^{crit} = 0.895$): a saddle, and trajectories drift slowly away along the nullclines. Middle: with a profit-led Kaleckian block the system is stable but close to divergence ($\sigma^{crit} = 0.937$). Right: wage-led, the nullclines cross, trajectories spiral clockwise into the steady state, and there is no divergence. Black lines are eigen-directions in the real-eigenvalue panels. The vertical scale of the third panel differs.

Economic interpretation. (i) *Same parameters, different fates.* The cost-only economy is unstable, the profit-led Kaleckian economy is stable but barely, and the wage-led economy is comfortably stable. The only difference is the sign and size of η_{net} . (ii) *In the profit-led case divergence is a creeping failure.* In the left and middle panels the system finds the nullclines quickly and then moves along them slowly; the nullclines are almost parallel and a growth shock leaves behind a long-lasting wage gap that keeps growth down. Adding the consumption channel moves the boundary out (from 0.895 to 0.937) because it takes away part of the competitiveness reward. (iii) *Wage-led: damped cycles.* Where the profit-led economy moves monotonically, the wage-led economy spirals. Its nullclines cross at an angle, with no direction along which a trajectory can run straight to or from the steady state, so every path is turned around it. The economic mechanism is the one described above: a boom raises productivity and lowers the wage share, which weakens consumption demand and ends the boom; the downturn then raises the wage share, and recovery follows. The cycle is damped because the loop is weak enough ($\sqrt{D} < 1$) that each round trip loses amplitude.

5.5 Stability in the two regimes

Proposition 2 (Distribution and stability). Profit-led ($\Lambda_{net} > 0$): *the eigenvalues are real and the system is stable iff $\sigma + \Lambda_{net} < 1$; monotonic divergence is still possible, but the threshold is higher*

than in the cost-only model. Wage-led ($\Lambda_{net} < 0$): all three Schur–Cohn conditions hold for every admissible parameter value, so there is no unstable region; the eigenvalues may be complex, with damped cycles of amplitude shrinking at rate $\sqrt{D} < 1$.

Table 2 gives the three conditions in the two regimes (derivations in Appendix A). In the profit-led row only $p(1)$ can fail; in the wage-led row none can.

	$p(1) > 0$	$p(-1) > 0$	$ D < 1$
Profit-led ($\Lambda_{net} > 0$)	$\phi[1 - \sigma - \Lambda_{net}]$. Can fail	$(1 + \tilde{\sigma})(2 - \phi) - \Lambda_{net}\zeta\phi$. Never fails: $> 2 - 2\phi$	$\sigma(1 - \phi) + \Lambda_{net}(1 - \zeta - \phi)$. Never fails when $p(1) > 0$
Wage-led ($\Lambda_{net} < 0$)	$\phi[1 - \sigma + \Lambda]$. Never fails	$(1 + \tilde{\sigma})(2 - \phi) + \Lambda \zeta\phi$. Never fails	$\sigma(1 - \phi) + \Lambda (\zeta + \phi - 1)$. Never fails

Table 2: The Schur–Cohn conditions in the two regimes ($|\Lambda| = -\Lambda_{net}$).




Figure 3: Which economies are stable? Each point is an economy defined by its net demand response η_{net} (negative: wage-led; positive: profit-led) and the strength σ of its accelerator–multiplier. A wage-led economy is stable; a profit-led economy is stable below the limit $\sigma + \Lambda_{net} = 1$ and diverges monotonically above it. Other parameters: $b = 0.5$, $\chi = 0.3$, $\zeta = 0.85$, $\phi = 0.2$.

The result is regime-contingent: the sign of Λ_{net} is the difference between a competitiveness channel and a consumption channel, which the model leaves free. It also says that the linear model

cannot, with either sign, generate a trap. What it can generate is monotonic divergence in a profit-led economy and damped cycles in a wage-led one, and nothing with a history: in both cases the economy ends up in the same place whatever its past. The next section adds what is missing.

6. A Nonlinear Model of Development Traps in a Dual Economy

6.1 Why the linear models are not enough

Sections 4 and 5 explain one mechanism well: productivity gains, relative costs and the distribution of income feed back on demand and growth. But as descriptions of a dual economy they miss what makes such economies puzzling. Ireland is the obvious case. Since 2008 the domestic economy has recovered strongly in some respects and remained fragile in others, and it is not obvious that a single shock, whether a banking crisis, the loss of an export market or a wave of foreign investment, leaves nothing behind once it has passed. This section asks what the linear models are missing and what economic ingredients would let two economies with the same fundamentals end up in very different places. A trap is a consequence of these ingredients, not the aim. Three features of the linear models sit badly with what dual economies look like:

1. **One economy, one destination.** A linear model has a single equilibrium. Given the parameters, the economy goes to the same place whatever its starting point, and a temporary shock, however large, leaves no mark once it has died away. Yet dual economies often look stuck, and a one-off episode (a banking crisis, the loss of a major export market, a wave of foreign investment) can leave a lasting imprint. Ireland's experience after 2008 is the obvious example.
2. **Instability with no brake.** When the profit-led loop is explosive, growth diverges without limit. Real economies do not do this. Something eventually stops the boom, and the form of the stopping point matters for what happens next.
3. **Distribution moves but never matters for the level of activity.** The wage share changes, but it feeds back only through the *growth rate* of productivity; its *level* never influences what firms do. In practice a squeezed profit share, or a fast-rising wage share, is exactly what changes the willingness to invest.

Each gap calls for a specific, defensible ingredient:

What the linear models lack	Economic ingredient	Where
Firms' behaviour is passive: no threshold	A kinked (threshold) investment function	Section 6.2
Growth that explodes, with no brake	A ceiling: external, supply or fiscal	Section 6.3
Distribution never feeds back in levels	The wage share as a slow state variable	Section 6.4

The argument in brief. The nonlinear model developed in this section rests on three ingredients, each motivated by the economics of a dual economy. First, firms invest only when demand growth is high enough to make expansion worthwhile, so low growth discourages the investment that would raise it. Second, growth is bounded above by a ceiling, whether an external, supply or fiscal limit, which binds even inside a currency union. Third, distribution adjusts slowly and feeds back on firms' willingness to invest, so the economy's history leaves a lasting mark.

A genuine trap, in the sense of Nelson and the poverty-trap literature, needs three things: a force that keeps the economy at low growth after small shocks, a force that carries it away from an intermediate threshold, and something that stops the escape at a high state that can be reached. The ingredients supply exactly these. The weak response of investment below the hurdle holds the economy at low growth, the strong response above it carries it away from the threshold, and the ceiling stops the escape. The slow-moving wage share makes the trap deepen with time. Together they allow two economies with identical fundamentals to settle at different long-run outcomes, one at low growth and one at the ceiling, depending on their history. All three are in the Harrod–Hicks–Goodwin tradition (Section 3). The system remains the one of Sections 4–5, with η_{net} in place of η throughout, and is not a separate model.

6.2 The kinked accelerator: economic motivation

In the base and Kaleckian models the only investment-like term is autonomous demand growth g_A , which responds to cost pressure but not to the state of demand itself. That leaves out the most basic Keynesian and Kaleckian fact about investment: firms add capacity when they expect demand to grow enough to use it, and when they can finance it. Three considerations shape the form we use.

A hurdle. A firm with spare capacity, facing sluggish demand, has no reason to expand. Investment responds to growth only once growth is high enough that existing capacity is likely to be fully used and the project pays. Below a hurdle growth rate q_F *induced* investment is zero; above it, faster growth calls forth more investment, which is the accelerator (Hicks, 1950). The function is therefore $\varsigma \max\{0, q_D - q_F\}$ and not a straight line. A straight line would make investment negative whenever growth fell below the hurdle, so that capital would be scrapped on a large scale. That is rare and slow: capital is long-lived and, once installed, cannot be sold off without loss. This irreversibility is the same reason that the investment literature gives for firms' waiting for a clear signal before committing (Dixit and Pindyck, 1994), and it is why a floor at zero is the realistic case. The hurdle also fits the stagnation tradition, in which investment is held back by excess capacity (Steindl, 1952).

Why the kink produces two different economies. Below the hurdle the economy has only the weak feedback of the base and Kaleckian models: each extra unit of growth adds little to next year's demand. Above it the accelerator adds a second, strong amplification on top (Figure 4b). The economy therefore has a *weakly self-reinforcing* regime at low growth and a *strongly self-reinforcing* one at high growth. The pieces fit together as follows. Weak feedback below the hurdle gives a resting place at low growth that attracts nearby paths. Strong feedback above it pushes paths away from an intermediate threshold. Once a ceiling is added, a second resting place appears at high growth. All three can coexist because the strength of amplification changes as growth crosses the hurdle. A linear function with the same strength everywhere cannot do this: it has a single equilibrium (Section 6).

What moves the hurdle. The hurdle is not a constant. It depends on profitability, hence on the wage share v_t (defined in Section 6.4): a higher wage share squeezes the profit margin and raises the hurdle ($\theta_v v_t$), which is the Kaleckian and Steindlian point that the profit share governs both the incentive and the ability (internal funds) to invest. It also depends on access to finance: a credit guarantee s makes marginal projects bankable and lowers the hurdle. These two terms connect the investment function to distribution and to policy, and the later sections use both.




Figure 4: The kinked accelerator. (a) Induced investment is zero until growth clears the hurdle q_F , then rises with slope ς . A higher wage share (orange) shifts the hurdle to the right; a credit guarantee (blue) shifts it to the left. (b) The share of this year's growth that carries into next year's demand growth: $1 - D_0 = 0.56$ below the hurdle (a shock fades), $1 - D_0 + k = 1.61$ above it (a shock is amplified), and nil once the ceiling binds. Calibration of Table 3.

A modelling choice, stated plainly

The $\max\{0, \cdot\}$ form is the simplest function with a hurdle and a floor, and it is piecewise linear, so every equilibrium and stability condition can be found by hand. A smooth S-shaped function such as a logistic would give the same qualitative picture. We choose the kinked form for tractability, not because firms literally behave this way.

The investment function. Firms invest only when growth clears the hurdle q_F , at which expansion is both profitable and financeable. The hurdle rises with the wage share v_t , which Section 6.4 introduces formally, and falls with the credit guarantee s :

$$I_t = \varsigma \max\{0, q_{D,t-1} - q_{F,t-1}\}, \quad q_{F,t} = q_F^0 + \theta_v v_t - s. \quad (12)$$

Induced investment enters demand with the domestic-content weight $(1 - \chi)$, so

$$q_{D,t}^d = \tilde{\sigma} q_{D,t-1} - \beta \Gamma_{t-1} + k \max\{0, q_{D,t-1} - q_{F,t-1}\} + \mu_{1,t}, \quad k \equiv (1 - \chi)\varsigma. \quad (13)$$

Below the hurdle (the *low regime*) the system is exactly that of Section 5, with own-feedback coefficient $\tilde{\sigma}$. Above it (the *ramp regime*) the coefficient is $\tilde{\sigma} + k$: growth begets more growth through investment. For either regime $p(1) = \phi [1 - A - \Lambda_{net} \zeta]$ with A the own-feedback coefficient, which gives

$$p(1)_{\text{low}} = \phi D_0, \quad p(1)_{\text{ramp}} = \phi (D_0 - k), \quad D_0 \equiv 1 - \sigma - \Lambda_{net}. \quad (14)$$

The low regime is stable whenever $D_0 > 0$. The ramp regime is a saddle exactly when

$$k > D_0 \iff \sigma + \Lambda_{net} + (1 - \chi)\zeta > 1. \quad (15)$$

This is a *local* Harrod knife-edge: growth is self-correcting below the hurdle and self-reinforcing above it. It is bounded below by the floor and, as the next step shows, above by the ceiling.

6.3 A ceiling on domestic growth

A strong accelerator would carry growth upward without bound if nothing stopped it. Realised growth is the smaller of what demand delivers and what the ceiling allows:

$$q_{D,t} = \min\{q_{D,t}^d, q_B\}. \quad (16)$$

The classical Thirlwall ceiling is $q_B = \varepsilon y_{ROW}/\pi$, export growth over the income elasticity of imports; if capital inflows finance a share λ_t of imports it becomes $q_B = \pi^{-1}[\varepsilon y_{ROW} + \dot{\lambda}/(1 - \lambda)]$, so that financing relaxes the constraint only while the financed share is rising (Thirlwall and Hussain, 1982). Inside a currency union the deficit country cannot run out of reserves, since TARGET2 balances accumulate automatically, so the constraint is soft; but it does not disappear. Two readings are in the literature. Sinn and Wollmershäuser (2012) stress the financing role of TARGET2; Cesaratto (2013) argues that it postpones but does not abolish balance-of-payments adjustment, since the substitute financing depends on decisions of the central bank and on the conditions attached. Ireland 2008–12 illustrates both: TARGET2 balances financed deposit flight, but the binding constraint was the bank–sovereign loop and the adjustment programme.

For Ireland the classical ceiling is the wrong tool for three reasons. The current account is in surplus and, in the headline figures, dominated by multinational activity; enclave exports do not respond to Irish prices or demand, so a ceiling computed for the whole economy would never bind; and what limits the *domestic* sector is leakage, capacity and fiscal and credit limits. We therefore treat q_B as the lowest of three ceilings on sector D :

$$q_B = \min\{q_B^{ext}, q_B^{sup}, q_B^{fis}\}, \quad (17)$$

an external ceiling (Thirlwall logic with domestic-sector elasticities, import intensity and profit repatriation), a supply ceiling (labour force, housing completions, infrastructure) and a fiscal ceiling (EU rules, sovereign and bank funding costs). The algebra needs only that q_B is a constant above the saddle growth rate; the interpretation determines how durable the high state is (Section 6.9).

6.4 The wage share as a third state, and the balanced-growth restriction

Because a higher wage share squeezes profits and raises the hurdle for investment, the model has to keep track of distribution as it evolves. Let v_t be the log deviation of the domestic wage share from a reference level. Since rulc_D is the growth of the wage share (Section 5),

$$v_t - v_{t-1} = \text{rulc}_{D,t} = \phi \Gamma_{t-1} - (1 - \zeta)(a + b q_{D,t-1}) + \psi h_t. \quad (18)$$

The wage share rises when catch-up and housing pressure outrun the part of productivity growth not passed into wages, and falls when productivity outpaces them.

In Part I's steady state the gap settles where domestic wage growth equals enclave wage growth: $w_D = \omega_F^0$, because the catch-up term $\phi \Gamma$ adjusts to deliver it. At any equilibrium of the growth-gap pair the wage share therefore drifts at

$$\text{rulc}^* = \omega_F^0 - a - b q^* = b(q^W - q^*), \quad q^W \equiv \frac{\omega_F^0 - a}{b}. \quad (19)$$

There are two conditions (a constant gap and a constant wage share) but only one free variable, growth q , which demand determines and which we cannot set. Both hold only at $q^* = q^W$. At any other growth rate the wage share keeps moving. With $\omega_F^0 = 2.4$, $a = 0.8$, $b = 0.5$ we have $q^W = 3.2$. In a low state with $q = 1.2$, productivity grows at 1.4 against wages at 2.4 and the wage share rises by one point a year; ... at a high state at the ceiling, $q = q_B = 3.6$ (Section 6.5), productivity grows at 2.6 and the wage share falls by 0.2 a year. A falling wage share is not a wage cut: wages still grow at 2.4. The economy thus has *fast* states (q_D, Γ) and a *slow* state v , and the drift of v will generate hysteresis.

6.5 The full system and an illustrative simulation

A note on the parameter values. The parameter values below are illustrative. They are chosen so that the model's mechanisms can be seen and so that the stability conditions of Sections 4–5 hold, and they are not estimated from, or matched to, Irish or other data. The numerical results that follow show which mechanisms matter and in which direction, not how large the effects are in practice.

$$\begin{aligned} q_{D,t} &= \min \left\{ \tilde{\sigma} q_{D,t-1} - \beta \Gamma_{t-1} \right. \\ &\quad \left. + (1 - \chi) \zeta \max \{ 0, q_{D,t-1} - q_F^0 - \theta_v v_{t-1} + s \} + \mu_{1,t}, q_B \right\}, \\ \Gamma_t &= -\gamma q_{D,t-1} + \delta \Gamma_{t-1} + \mu_{2,t}, \\ v_t &= v_{t-1} + \phi \Gamma_{t-1} - (1 - \zeta)(a + b q_{D,t-1}) + \psi h_t. \end{aligned} \quad (20)$$

Public autonomous demand growth g_G enters $\mu_{1,t}$ with weight $(1 - \chi)$. The wage share affects the fast pair only through the hurdle q_F . The numbers below are chosen for illustration and are not estimates; they keep $\sigma + \Lambda_{net} < 1$, so Sections 4–5 are stable on their own.

Parameter	Value	Parameter	Value
b	0.5	σ	0.5
χ	0.3	ς	1.5
ζ	0.85	ℓy_F	0.726
ϕ	0.2	g_A^0, g_G, s	0, 0, 0
η, s_C, ξ_0	0.3, 0.2, 0.3	$a, \omega_F^0, \psi h$	0.8, 2.4, 0.2
η_{net}	0.18 (profit-led)	q_F^0, θ_v, q_B	2.2, 0.02, 3.6

Table 3: Illustrative simulation.

Derived values: $\Lambda_{net} = 0.063$, $\tilde{\sigma} = 0.509$, $\beta = 0.0252$, $\gamma = 0.425$, $\delta = 0.8$, $D_0 = 0.437$, $k = 1.05$, $\mu_1 = 0.716$, $\mu_2 = 1.52$, $q^W = 3.2$.

6.6 Three equilibria

Fast subsystem. The wage share moves slowly, so we first hold v fixed, which fixes q_F , and ask where (q_D, Γ) comes to rest. Every equilibrium has Γ on its nullcline

$$\Gamma = \frac{\mu_2 - \gamma q}{\phi}, \quad (21)$$

which slopes down: faster growth narrows the gap. Two equations would give one resting point in each linear regime; what gives three is that the economy behaves in three different ways depending on how fast it grows (no investment below the hurdle, accelerating investment above it, the ceiling). A candidate equilibrium is valid only if it lies in the regime in which its own equations apply. Solving each regime's equations on (21) gives

$$q_L = \frac{N_0}{D_0}, \quad q_S = \frac{k q_F - N_0}{k - D_0}, \quad q_H = q_B, \quad N_0 \equiv \mu_1 - (1 - \chi) \eta_{net} \mu_2, \quad (22)$$

with $\Gamma_i = (\mu_2 - \gamma q_i)/\phi$. Moreover $q_S - q_F = D_0(q_L - q_F)/(D_0 - k)$, so when $k > D_0$ and $q_L < q_F$ the saddle lies above the hurdle: it exists exactly when the low state does.

Proposition 3 (Three equilibria). *Suppose $k > D_0$. The fast subsystem has three equilibria L, S, H if and only if $q_L < q_F$ and $q_B > q_S$. If $q_L \geq q_F$ the trap and the saddle vanish and only H remains; if $q_B \leq q_S$ only the trap remains. If $k \leq D_0$ there is a single stable equilibrium.*

With the calibration, $q_L = 1.20$ ($\Gamma_L = 5.05$), $q_S = 2.913$ ($\Gamma_S = 1.41$) and $q_H = 3.6$ ($\Gamma_H = -0.05$). The economics of each:

- **L , the trap.** Growth of 1.2 per cent lies below the 2.2 at which firms invest, so investment is zero; demand growth rests on the enclave linkage and the weak multiplier. Slow growth means little productivity gain, so the wage gap stays wide. Nothing in the state generates a push beyond the hurdle.
- **S , the tipping point.** The growth rate at which investment becomes self-reinforcing. It lies above the hurdle (2.91 against 2.2) because leakage and wage-catch-up drag take back part of each point of growth. It is a threshold, not a place the economy can stay, and exists only because L does.
- **H , the ceiling state.** Demand would deliver 4.02 but growth is held at the ceiling 3.6. The wage gap has closed ($\Gamma_H \approx 0$). It is a constrained, demand-led state, not one of unlimited growth.

6.7 Stability

6.8 Stability

Near each equilibrium the map is linear. Holding the wage share fixed, the Jacobian of the fast pair (q_D, Γ) in each regime is

$$J_{\text{low}} = \begin{pmatrix} \tilde{\sigma} & -\beta \\ -\gamma & \delta \end{pmatrix}, \quad J_{\text{ramp}} = \begin{pmatrix} \tilde{\sigma} + k & -\beta \\ -\gamma & \delta \end{pmatrix}, \quad J_{\text{ceiling}} = \begin{pmatrix} 0 & 0 \\ -\gamma & \delta \end{pmatrix}. \quad (23)$$

Equivalently, $J = \begin{pmatrix} A & -\beta \\ -\gamma & \delta \end{pmatrix}$ with $A = \tilde{\sigma}$ in the low regime, $\tilde{\sigma} + k$ in the ramp regime and 0 in the ceiling regime, where the top-right entry is also 0. At the ceiling, growth is set by q_B and depends neither on its own past nor on the gap, so the whole first row vanishes.

Numerical illustration. Table 4 gives the eigenvalues and type of each equilibrium for the illustrative calibration of Table 3, with $b = 0.5$, $\chi = 0.3$, $\zeta = 0.85$, $\phi = 0.2$, $\eta_{\text{net}} = 0.18$, $\sigma = 0.5$, $\varsigma = 1.5$ and $q_B = 3.6$. The whole difference between L and S is the single term k in the top-left entry.

Equilibrium	q_D	Γ	Eigenvalues	Type
L (trap)	1.20	5.05	0.476, 0.833	stable node
S (tipping point)	2.91	1.41	1.573, 0.786	saddle
H (ceiling)	3.60	-0.05	0, 0.8	stable

Table 4: Three equilibria: trap, tipping point and ceiling. Eigenvalues of the regime Jacobian at $v = 0$ ($q_F = 2.2$).

L is stable ($p(1) = \phi D_0 = 0.087 > 0$, real eigenvalues): small shocks lift growth temporarily but not above the hurdle, so it drifts back. S is a saddle ($p(1) = \phi(D_0 - k) = -0.123 < 0$): above it growth feeds on itself, below it growth returns to L . H is stable because the ceiling decouples growth from its own past, so the only moving variable is the gap, closing at speed ϕ ; it is only as durable as the ceiling.

Phase diagrams. Figure 5 shows the one-dimensional picture and Figure 6 the phase diagram. The separatrix, the stable manifold of S , is almost vertical because β is small: whether an economy escapes depends mainly on its growth, not on its wage gap. At L the push needed is about $q_S - q_L = 1.7$ points of growth.




Figure 5: Demand growth along $\Delta\Gamma = 0$ against the 45° line. Where the curve lies above the line growth rises. Below the hurdle the slope is below one (stable crossing L); above it the slope exceeds one (unstable crossing S); the curve flattens at the ceiling, giving the stable crossing H . The dashed curve shows a credit guarantee $s = 0.6$ shifting the kink left and moving S from 2.91 to 1.89.




Figure 6: Phase diagram of the fast subsystem at $v = 0$. Blue: $\Delta q = 0$. Orange: $\Delta \Gamma = 0$. The near-vertical line through the saddle S is the separatrix: starts to its left converge to L (blue shading), starts to its right to H (orange shading). Grey lines are sample trajectories.

6.9 Hysteresis: the wage share moves the barrier

Now let v move. The hurdle is $q_F(v) = q_F^0 + \theta_v v - s$, so the barrier is

$$q_S(v) = \frac{k q_F(v) - D_0 q_L}{k - D_0}, \quad \frac{\partial q_S}{\partial v} = \theta_v \frac{k}{k - D_0} = 0.034. \quad (24)$$

The barrier rises faster than the hurdle (factor 1.71). The story has three steps. In the trap wages keep following the enclave while productivity is slow, so the wage share rises by one point a year. A rising wage share squeezes profits and raises the hurdle, so the barrier to escape rises. Each year in the trap therefore makes escape harder. At the ceiling the reverse holds: productivity outpaces wages, the wage share falls and the high state is reinforced. The high state exists while $q_B > q_S(v)$, i.e. for

$$v < \hat{v}_{hi} = \frac{1}{\theta_v} \left[\frac{q_B(k - D_0) + D_0 q_L}{k} - q_F^0 + s \right] = 20.1. \quad (25)$$

In the trap the wage share rises by one point a year, so after about twenty years the economy reaches $\hat{v}_{hi}$ and the high state disappears altogether. Figure 7 shows the resulting bifurcations. Two economies with identical technology, trade structure and policy can be in different positions simply because one has spent twenty years in the trap. This is the dual-economy analogue of scarring in hysteresis arguments: low-growth episodes do not merely delay recovery, they raise its cost, and

the mechanism is distributional.




Figure 7: How policy and the wage share move the economy's resting places. Solid lines: places where the economy can settle; dashed: the tipping point. (a) A stronger credit guarantee lowers the growth firms need before they invest, so the tipping point falls; once low enough the ceiling state is reachable, and with a stronger guarantee still the trap disappears. (b) A rising wage share does the reverse: the tipping point climbs, and when it reaches the ceiling the high state is lost.

Durability of the high state. At H the wage share drifts at $\text{rulc}_H = b(q^W - q_B)$. If $q_B \geq q^W$ the wage share does not rise and the high state is self-sustaining. If $q_B < q^W$ the wage share rises slowly, the barrier rises with it, and when $q_S(v)$ reaches q_B the high state disappears. With $q_B = 3.0 < q^W = 3.2$ the high state lasts about 25 years and the economy then returns to the trap. A ceiling below the wage-neutral growth rate is one at which wages keep catching up with productivity until profitability erodes; a high state is durable only if growth at the ceiling is fast enough to keep productivity ahead of wages. This is a result of the model, not an assumption, and it ties the interpretation of the ceiling (external, supply or fiscal) to the durability of recovery.

6.10 Is the trap a profit-led phenomenon?

The illustration is profit-led ($\eta_{\text{net}} = 0.18$), so one may ask whether the trap needs it. It does not. The trap exists whenever (i) $k > D_0$, (ii) $q_L < q_F$ and (iii) $q_B > q_S$, and these can hold in either regime (Appendix B). In a wage-led economy D_0 is larger, so a stronger accelerator is needed, and with the same autonomous demand the low state sits higher and the trap is shallower; in a strongly wage-led economy the saddle can lie above the ceiling, so there is no high state to escape to. The wage share rises in the trap in both regimes, so hysteresis does not depend on the regime. In a

profit-led economy growth feeds itself, which the accelerator magnifies into a tipping point; in a wage-led economy faster growth lowers the wage share and cuts the consumption that was holding demand up, so the same accelerator must overcome a drag. A wage-led trap is a stagnation of investment, not consumption. A caveat: the wage share enters only through the investment hurdle and not through the level of consumption, so the wage-led results are a first approximation; the Goodwin step (Section 11) is the proper treatment.

6.11 Two extensions: capacity utilization and employment

Two questions arise. The model has no role for capacity utilization, and its wage share drifts without bound. We tested one extra state for each, using simple specifications chosen for illustration; the results are indicative.

Capacity utilization. Add a state u (the deviation of utilization from normal) with capacity growth $\bar{g} + \kappa u$, utilization changing by $q - g_K$, and let idle capacity raise the hurdle with strength ρ_u . No new equilibrium appears and the system gains one eigenvalue, $1 - \kappa$. The barrier shrinks as the hurdle becomes more responsive to utilization (from 3.04 to 2.55 to 2.37 for $\rho_u = 0, 0.3, 0.6$) and rises with slack in capacity growth. With $\kappa = 0$ the Harrodian knife-edge returns and the supply ceiling can be read as maximum utilization.

Employment (the Goodwin step). Add an employment state that changes by growth minus productivity growth minus labour-force growth, and let wage pressure rise with employment. If the pressure enters *domestic* wages nothing changes: the gap Γ absorbs it, as with cost relief (Section 8). If it enters the *enclave* wage anchor, because the two sectors share a labour and housing market, the wage share no longer falls automatically at the high state: with a moderate response the high state is lost after decades, and with a strong one one-way hysteresis becomes a slow stop-go cycle. This is the starting point for the Goodwin extension.

7. Robustness: What Depends on Dualism?

A fair criticism is that the dual structure is put in by hand, through the investment hurdle and the weak linkages, rather than emerging from the dynamics. What is *assumed* is a maintained structural fact: an enclave whose growth does not depend on the domestic sector and which is weakly linked to it. In Ireland it is the historical result of multinational entry and the tax regime, and the model does not explain its origin. What the model *generates* is whether the structure persists: the low-growth state, the tipping point, the ceiling state and the hysteresis. Three tests show which is which.

Test 1: remove the enclave from the hurdle. Replace the enclave-linked hurdle by a constant ($\theta_v = 0$), and separately switch off the wage-gap channel ($\eta_{\text{net}} = 0$). In both cases the three equilibria survive: $q_L = 1.20$ and $q_S = 2.91$ with a constant hurdle, and $q_L = 1.45$ and $q_S = 2.88$ without the wage-gap channel. The trap therefore does not need dualism. It comes from the floor, the kinked accelerator and the ceiling. Dualism determines how deep the low state lies below the hurdle, and the hysteresis, which needs the hurdle to move with the wage share. It does not create the trap; it determines whether the economy sits in it.

Test 2: let the dualism parameters vary. Two parameters set how dual the economy is: the strength of the enclave's pull on domestic demand (ℓy_F) and the leakage χ . Figure 8 varies both, with everything else at the calibration. With stronger linkages q_L rises and crosses the hurdle at $\ell y_F \approx 1.16$ (against 0.73 at the calibration), beyond which the low state is gone. With higher leakage the accelerator weakens, the tipping point rises above the ceiling, and the trap becomes a lock-in with no high state. Dualism is not a switch but a region of parameter space: the trap is the middle case, and a well-linked economy is not dual at all.




Figure 8: Dualism as a region of parameter space. Horizontal axis: how strongly the enclave pulls up domestic demand. Vertical axis: how much of each extra unit of domestic demand leaks out. Blue: the trap, with a low and a high state both stable. Brown: only the low state exists. Yellow: the enclave pull lifts domestic growth above its hurdle and the dualism disappears. The dot is the calibration.

Test 3: the wage anchor. The model assumes that domestic wages are pulled towards the enclave level (the catch-up term $\phi\Gamma$) while enclave wage growth is set outside the model. The evidence for the first is indirect. Varying the catch-up speed ϕ from 0.02 to 0.8 leaves the trap in place, because the trap conditions do not contain ϕ ; ϕ only scales the equilibrium wage gap. Varying enclave wage growth ω_F^0 from 0 to 4 also leaves it in place, even with no enclave anchor, but a higher enclave wage norm moves the low state down (1.89 to 0.74) and the tipping point up (2.42 to 3.24). The assumption therefore matters for how deep the trap is, for the wage-share drift behind the hysteresis, and for the results on wage-level shocks and internal devaluation (Section 9), which are conditional on it.

What remains. The model treats linkages and leakage as parameters. If they eroded when the domestic sector lags the enclave, a well-linked economy could drift into the trap and dualism would become an attractor. A full account of how an enclave arises is a different project (Section 11).

8. Policy: Fiscal Levers and Timing

A fiscal lever here is an instrument of the state that changes the conditions under which firms and households act and so shifts the equilibria and the tipping point. Ireland has no national interest rate or exchange rate, so the macro instruments it controls are public spending, taxation and state-backed credit. Three levers map onto three channels of the model: the investment hurdle, demand and costs. They are treated as permanent parameter shifts (or, in the timing experiment, as temporary pulses), and the model does not price them; the only fiscal limit is the ceiling q_B . We ask of each whether it removes the trap ($q_L \geq q_F$), whether it shrinks the barrier $q_S - q_L$, and whether it matters at all.

Credit guarantee s . It enters only through the hurdle $q_F = q_F^0 + \theta_v v - s$, and $\partial q_S / \partial s = -k / (k - D_0) = -1.71$. The trap disappears when the hurdle falls to the low state's growth rate, at $s^* = q_F^0 + \theta_v v - q_L = 1.00 + 0.02v$: at that point the saddle and the trap collide and vanish (a border-collision bifurcation). Below s^* the trap persists but a smaller push suffices to escape. The guarantee works on firms' willingness to commit: it lowers the lender's risk, so projects that were not bankable at the low state's growth rate become so. The cost is borne by the state as a contingent liability, and s^* rises by $\theta_v = 0.02$ per point of wage share, so delay raises the guarantee needed.

Public autonomous demand g_G . It enters μ_1 with weight $(1 - \chi)$, so $\partial q_L / \partial g_G = (1 - \chi) / D_0 = 1.60$ and $\partial q_S / \partial g_G = -(1 - \chi) / (k - D_0) = -1.14$: the barrier shrinks by 2.74 per point. The trap vanishes at $g_G^* = (q_F - q_L) D_0 / (1 - \chi) = 0.62$ percentage points. Public demand does not change

firms' beliefs but adds demand directly, and because it is autonomous it keeps demand rising when private investment is dormant, amplified by the multiplier $1/D_0$. But g_G is a *growth rate*: removing the trap requires public demand to grow 0.62 points faster year after year, a standing commitment that runs into the fiscal ceiling.

Cost relief (housing costs). A fall in ψh shifts both forcing terms: through μ_1 by $-(1-\chi)\eta_{net}\psi h$ and through μ_2 by $+(1-\chi)\eta_{net}\psi h$. They cancel exactly, so q_L , q_S and the wage-share drift are unchanged and only the level of the gap Γ shifts. In equilibrium wage growth is anchored at the enclave rate, so housing costs only change how much catch-up pressure is needed to deliver it. Cost relief is the weakest lever for the *growth trap*. This is a statement about this structure, not about the policy's value: relief that eases the supply ceiling raises q_B , which matters for how long the high state lasts.

Lever	Effect on $q_S - q_L$	Removes the trap at	Channel
Credit guarantee s	-1.71 per unit	$s^* = 1.00 + 0.02v$	firms' willingness to invest
Public demand g_G	-2.74 per pp	$g_G^* = 0.62$ pp	demand, via multiplier $1/D_0$
Cost relief ψh	none	never	shifts Γ only

Table 5: Comparison of levers at $v = 0$.

Two levers can end the trap, one cannot. The guarantee lowers the growth rate firms need to see before they invest; public demand raises the growth the economy delivers while investment is dormant. The trap ends when the two meet. A guarantee is a one-off change of level and its cost arises only if projects fail; public demand must grow faster than output year after year, so it is limited sooner by the fiscal ceiling, while the guarantee does nothing when there is no demand to invest for. At the margin one point of public-demand growth does the work of about 1.6 units of guarantee, and in practice they are complements.

Timing. While the economy sits in the trap the wage share rises, the hurdle rises with it, and the same policy has a different effect at a different date. Figure 9 applies the same temporary public-demand push (two points above trend for three years) early and late to an economy starting in the trap. Early (year 5), growth jumps to the ceiling and stays there; once there the wage share starts to fall because at the ceiling productivity grows faster than wages. Late (year 30), growth reaches the ceiling for a few years and falls back: the wage share has passed the point of no return and no high state exists. A permanent credit guarantee, which brings the hurdle back to its starting value, keeps the late push at the ceiling. In the calibration the latest start at which the two-point push still works is year 13, and the smallest push that works at year 5 is 1.5 points. A temporary push is a

kick: useful when there is a high state to be kicked towards. When the window has closed only a lever that reshapes the landscape works.




Figure 9: The same policy applied early and late to an economy starting in the trap. Public demand growth is raised by 2 points for three years. Blue (year 5): growth reaches the ceiling and stays. Orange (year 30): growth reaches the ceiling briefly, then falls back because the wage share has pushed the hurdle too high. Dashed orange: the late push with a permanent credit guarantee. Grey: no policy.

Escape is only half the problem; staying out is the other. The high state is durable only if growth at the ceiling is at least q^W , so supply-side policy that raises q_B does not help the economy escape but helps the escape last. Demand-side policy gets the economy out; supply-side policy keeps it out.

8.1 Fiscal Rules and the Enclave Cycle

The levers above are chosen once. In practice public demand also responds, by rule or by habit, to the enclave's fortunes, and the way it responds decides whether an enclave bust costs the economy its high state. Let w_t be the enclave demand shock (the movement in the foreign-linked term ℓ_{yF}) and let public demand growth follow $g_{G,t} = -\kappa w_t$. Three rules span the cases:

- *procyclical* ($\kappa = -1$): spending rises with receipts in a boom and is cut in a bust, as when windfall corporation tax is spent as it arrives and then withdrawn (Kaminsky, Reinhart and Végh 2004; Tornell and Lane 1999);
- *flat* ($\kappa = 0$): public demand growth ignores the enclave;
- *countercyclical* ($\kappa > 0$): windfalls are saved and spent in the bust, as under a savings rule that feeds a fund such as the Future Ireland Fund.

The enclave shock enters demand with weight one and public demand with weight $1 - \chi$, because a share χ leaks to imports. The net shock is therefore $(1 - \kappa(1 - \chi))w_t$, and a rule with $\kappa = 1/(1 - \chi) = 1.43$ cancels it exactly.

Proposition 4. *Starting in the high state, a permanent enclave demand loss $|w|$ leaves the high state in place if and only if*

$$(1 - \kappa(1 - \chi))|w| < w^*(v) = (q_B - q_S(v))(k - D_0). \quad (26)$$

Proof. The net shock raises the saddle by $(1 - \kappa(1 - \chi))|w|/(k - D_0)$, and the high state survives while the saddle stays below q_B . $\square$

In the calibration $w^*(0) = 0.42$ points, and it falls by 0.02 per point of wage share. A flat rule tolerates $|w| < 0.42$. A procyclical rule multiplies the shock by 1.7 and tolerates only $|w| < 0.25$, so a 0.3 bust that a flat rule absorbs is fatal to it. A countercyclical rule with $\kappa = 1$ tolerates $|w| < 1.40$, and one with $\kappa = 1/(1 - \chi)$ tolerates any loss, at the price of spending $1.43|w|$ more every year the loss lasts. Saving in a boom is costless: in the high state demand runs above the ceiling (4.2 to 4.6 in the simulations), so extra spending adds nothing to growth and withheld spending costs nothing. This is the model's case for the savings rule.

Temporary busts and hysteresis. Figure 10 applies a boom of 0.3 for five years followed by a bust of 0.6 for five years, after which the shock vanishes. Under the flat rule, growth falls during the bust, passes below the saddle and, once the shock is gone, settles in the *low* state (1.2 per cent): a temporary loss has a permanent effect. The procyclical rule does the same, faster and deeper. The countercyclical rule keeps growth at 3.6 throughout, with a net outlay equal to the bust minus the boom. Panel (c) shows the wage share falling at the ceiling and rising, with the hurdle, in the two economies that fall into the low state.




Figure 10: Three fiscal rules under an enclave boom (years 0–4, $+0.3$) and a temporary bust (years 5–9, -0.6), starting in the high state. Dashed grey: procyclical ($g_G = w$). Blue: flat ($g_G = 0$). Orange: countercyclical ($g_G = -w/(1 - \chi)$). Both other rules fall below the saddle during the bust and end in the low state.

Limits. The countercyclical rule stabilises *transitory* shocks. For a permanent loss it requires g_G to stay $1.43|w|$ above trend indefinitely, which runs into the fiscal ceiling: in a permanent 0.3 downturn the cumulative extra public demand reaches 12.9 points of growth over forty years. It also requires telling transitory shocks from permanent ones, which cannot be done in real time. The model has no fiscal block, so it says nothing about the size of the fund needed or the financing of the bust.

9. Application: Ireland After 2008

This section applies the model to Ireland after 2008, when the state absorbed large bank debts and responded with fiscal consolidation and an internal devaluation, that is, an attempt to restore competitiveness by cutting unit labour costs. This is a stylised illustration, not an estimate. The numbers use the illustrative parameter values of Table 3, which are chosen to show which mechanisms matter and in which direction, and are not matched to Irish data.

Feature of the episode	Model object	Shock used
Fiscal consolidation	public demand growth g_G	g_G down 0.5 pp, permanent
Banking crisis, credit crunch	credit guarantee s	$s = -0.2$, permanent
Bank debt lowers the fiscal limit	ceiling q_B	q_B from 3.6 to 2.6
Internal devaluation	wage cost term	wage growth down 2 pp for three years
Enclave export boom	linkage term ℓ_{y_F}	up by 1.0
ECB backstop	ceiling restored	$q_B = 3.6$ from year 9

Table 6: Mapping the episode onto the model. The economy starts at H ; the crisis hits in year 5.

The package changes the equilibria as follows. Before the crisis $q_L = 1.20$, the hurdle is 2.20 and the barrier is 2.91. After austerity and the credit crunch $q_L = 0.40$, the hurdle 2.40 and the barrier 3.83; since the lowered ceiling (2.6) is below the barrier, no high state exists. After the enclave boom $q_L = 2.69$ exceeds the hurdle and the trap itself disappears.




Figure 11: Stylised crisis paths from the high state (crisis in year 5). Orange: austerity, credit crunch and lower ceiling. Dashed: the same plus a three-year wage cut. Blue: the same plus an enclave export boom and a restored ceiling from year 9.

Three results stand out. First, *austerity alone sends the economy to the trap and keeps it there*: growth falls to the lowered ceiling and then drifts to $q_L = 0.4$, and the wage share, which had been falling in the boom, rises about a point a year. Second, *internal devaluation helps briefly and then fades*: a three-year wage cut lifts growth by about 0.3 points in years 7–8 and lowers the wage

share by about five points at the trough, but the lift is gone by year 10 and growth converges to the same trap, because a change in the wage *level* shifts Γ and v but not q_L or q_S . Third, *recovery comes from outside the domestic loop*: the enclave boom raises ℓ_{y_F} , lifts q_L above the hurdle and removes the trap; with the ceiling restored, growth reaches q_B within about seven years of the shock.

Three features work against internal devaluation in this structure. Domestic wage growth is anchored by the enclave, so a cut in the level fades as wages catch up. Productivity is demand-driven (Verdoorn), so contraction lowers productivity growth and offsets the cut in unit labour costs; productivity gains observed after 2008 were partly composition effects, as low-productivity jobs in construction disappeared. And whether a lower wage share helps at all depends on whether the economy is profit-led or wage-led; in a wage-led economy the cut would deepen the slump. The ceiling matters too: bank debt enters through the fiscal limit, a low q_B not only caps growth but can eliminate the high state, and the central-bank backstop is what relaxes it, which connects the episode to the TARGET2 discussion of Section 6. That an enclave export boom carried the recovery is a hypothesis for Ireland, not a result: it needs to be checked against data on enclave exports, the linkage to domestic demand (the GDP–GNI* gap is relevant) and the timing of the backstop. The conclusion is not that wage restraint is harmless or helpful in general; it is that in this structure it cannot by itself change which equilibrium the economy ends up in.

10. Relation to Thirlwall’s Law

It is worth being precise about what has and has not been generalised. Thirlwall’s law states $g = x/\pi$: growth is constrained by export growth divided by the income elasticity of import demand, because the constraint is ultimately a foreign-exchange or financing constraint on the whole economy. Three modifications were needed to make it relevant to Ireland.

First, the constraint is moved from the whole economy to the domestic sector considered separately from the enclave, since the enclave’s export earnings, not the domestic sector’s, dominate the aggregate current account, and the whole economy is not constrained: it runs a persistent external surplus (Irish Fiscal Advisory Council, 2023). The “export” analogue in equation (1) is the enclave’s internal transfer to the domestic economy, ℓ_{y_F} .

Second, the source of the constraint is relocated from a price and quantity trade elasticity to a linkage and leakage structure. The limiting factor is not Ireland’s ability to earn foreign exchange but the domestic economy’s ability to capture, via wages, taxation and procurement, a share of value added generated within its borders by firms whose production decisions are exogenous to it. This is arguably a *tighter* constraint than the classical one, because ℓ is not a price-adjustable elasticity: there is no devaluation that increases local procurement by a pharmaceutical multinational’s Irish subsidiary.

Third, the external constraint is not abandoned but demoted to one component, q_B^{ext} , of a composite ceiling on the domestic sector (Section 6). Inside a currency union the classical financing constraint is soft, though not absent, and supply and fiscal limits may bind first.

The cumulative-causation structure addresses a feature Thirlwall flagged but did not formalise: that a country's specialisation pattern is not fixed but can be reinforced by the growth process the constraint generates (Thirlwall, 2011). In the Irish case that reinforcement runs through the labour market and the wage share, not through trade structure.

11. Limitations and Future Work

- **Dualism is assumed at the level of origin.** Section 7 shows what does and does not depend on it. Making linkages and leakage state variables that erode when the domestic sector lags the enclave would turn dualism into an attractor; modelling enclave formation (increasing returns, core-periphery) is a different project.
- **The wage anchor.** Domestic wages follow enclave wage growth through catch-up. The trap does not need this, but its depth, the hysteresis and the weakness of internal devaluation do. Letting the anchor respond to domestic conditions, such as labour-market tightness or the enclave's share of employment, is future work.
- **The wage share enters only through the investment hurdle,** not the level of consumption, and its drift is unbounded in both states, so the fast-subsystem equilibria are full equilibria only if $q^* = q^W$. Closing the loop with wage growth responding to employment, a Goodwin predator-prey structure with an enclave wage anchor, is the natural next step, and the setting in which wage-led and profit-led regimes can be treated fully.
- **No utilization or financial block.** Section 6.11 shows one extra state each for utilization and employment changes the barrier but not the logic. Financial constraints are summarised in the guarantee s and the ceiling q_B .
- **Illustrative calibration.** All numbers show direction and mechanism, not magnitudes for Ireland.

12. Conclusion

This paper asked whether Ireland's dual structure, a foreign-owned enclave whose growth is exogenous to the domestic economy and only weakly linked to it, can trap the *domestic* economy at low growth, and through which mechanism. We took Thirlwall's balance-of-payments-constrained growth as the natural starting point and asked what survives of it inside a currency union with a

persistent external surplus. We then built a linkage-constrained model of the domestic economy and used it to study traps, hysteresis and fiscal policy.

The first finding is negative. Ireland's dual economy does not generate a classical Thirlwall trap. The country runs a persistent external surplus, has no exchange-rate margin as a currency-union member, and its dominant export sector is unresponsive to the price and income variables that discipline the classical constraint. A literal application of $g = x/\pi$ to Irish trade data would estimate an artefact of contract manufacturing and intellectual-property relocation.

There is nonetheless a genuine developmental trap, and it runs through the labour market and the weakness of linkages between the enclave and the domestic economy.

A linkage-constrained growth condition, combined with Verdoorn's law and wage catch-up, gives a two-state system whose only instability is monotonic divergence at a Harrodian knife-edge, which a Kaleckian consumption channel removes in a wage-led economy. A trap proper needs two things: firms that invest little when growth is low and strongly when it is high, and a ceiling that stops growth from rising without limit. Together they allow an economy to rest at low growth, at high growth, or at an unstable threshold between them. Distribution adds history. While the economy stays in the trap, productivity grows more slowly than wages, so the wage share rises, profits are squeezed and firms need stronger demand before they will invest. Each year in the trap therefore makes escape harder, so the same policy works early and fails late. The high state, in turn, lasts only if growth at the ceiling is fast enough to keep productivity ahead of wages. A credit guarantee and public demand can end the trap, cost relief cannot, and internal devaluation changes the path but not the destination.

The result survives removing dualism from the hurdle, switching off the wage-gap channel, varying the speed of wage catch-up and the enclave wage norm, and moving between profit-led and wage-led regimes, though the depth of the trap and the hysteresis depend on them. The natural next steps are the Goodwin extension, endogenous linkages and estimation of the parameters, particularly the linkage coefficient, the regime sign η_{net} and the investment hurdle, to establish whether Ireland's domestic economy is currently below or above the tipping point.

A. Derivations

Reduced form. Substitute Verdoorn's law into relative unit labour cost: $\text{rulc}_{D,t} = \phi \Gamma_{t-1} - (1 - \zeta)a - (1 - \zeta)bq_{D,t-1} + \psi h_t$. Put this into the feedback equation and multiply by $(1 - \chi)$:

$$(1 - \chi)g_{A,t} = (1 - \chi)[g_A^0 + \eta(1 - \zeta)a - \eta\psi h_t] + (1 - \chi)\eta(1 - \zeta)bq_{D,t-1} - (1 - \chi)\eta\phi\Gamma_{t-1}.$$

Inserting into (1), the coefficient on $q_{D,t-1}$ is $\sigma + (1 - \zeta)\Lambda \equiv \tilde{\sigma}$, on Γ_{t-1} is $-(1 - \chi)\eta\phi \equiv -\beta$, and the remainder is $\mu_{1,t}$. For the gap, $\Gamma_t = \Gamma_{t-1} + \omega_F^0 - \zeta(a + bq_{D,t-1}) - \phi\Gamma_{t-1} - \psi h_t$, so the coefficient on Γ_{t-1} is $1 - \phi = \delta$, on $q_{D,t-1}$ is $-\zeta b = -\gamma$, and $\mu_{2,t} = \omega_F^0 - \zeta a - \psi h_t$.

Schur–Cohn conditions. Let A be the own-feedback coefficient. Then $T = A + \delta$, $D = A\delta - \beta\gamma$, and $\beta\gamma = (1 - \chi)\eta\phi\zeta b = \Lambda\zeta\phi$. Hence

$$p(1) = 1 - T + D = (1 - A)(1 - \delta) - \beta\gamma = \phi[1 - A - \Lambda\zeta].$$

For $A = \tilde{\sigma} = \sigma + (1 - \zeta)\Lambda$ this equals $\phi[1 - \sigma - (1 - \zeta)\Lambda - \zeta\Lambda] = \phi[1 - \sigma - \Lambda]$, so ζ cancels. Next, $p(-1) = (1 + \tilde{\sigma})(2 - \phi) - \Lambda\zeta\phi > 2 - 2\phi > 0$ because $\tilde{\sigma} \geq 0$ and $\Lambda\zeta\phi < \phi$. If $p(1) > 0$ then $\tilde{\sigma} \leq \sigma + \Lambda < 1$, so $D = \tilde{\sigma}\delta - \Lambda\zeta\phi < 1$, and $D \geq -\Lambda\zeta\phi > -\phi > -1$. The discriminant is $T^2 - 4D = (\tilde{\sigma} - \delta)^2 + 4\Lambda\zeta\phi$.

With $\Lambda_{net} < 0$, write $|\Lambda| = -\Lambda_{net}$; then $|\eta_{net}| \leq s_C \xi_0 < 1$ so $|\Lambda| < (1 - \chi)b < 1$, $p(1) = \phi[1 - \sigma + |\Lambda|] > 0$ and $p(-1) = (1 + \tilde{\sigma})(2 - \phi) + |\Lambda|\zeta\phi > 0$ since $\tilde{\sigma} > -1$. Further $D = \sigma(1 - \phi) + |\Lambda|(\zeta + \phi - 1)$: if $\zeta + \phi \leq 1$ then $D \leq \sigma(1 - \phi) < 1$; if $\zeta + \phi > 1$ then $|\Lambda|(\zeta + \phi - 1) < \phi$, so $D < \sigma + \phi(1 - \sigma) < 1$. The discriminant becomes $(\tilde{\sigma} - \delta)^2 - 4|\Lambda|\zeta\phi$, which is negative when $\tilde{\sigma}$ is close to $\delta = 1 - \phi$, giving a conjugate pair with $|\lambda| = \sqrt{D} < 1$.

Equilibria. In the low regime put (21) into $q = \tilde{\sigma}q - \beta\Gamma + \mu_1$: $q(1 - \tilde{\sigma} - \beta\gamma/\phi) = \mu_1 - (\beta/\phi)\mu_2$. Using $\beta\gamma/\phi = \Lambda\zeta$, $1 - \tilde{\sigma} - \Lambda\zeta = D_0$ and $\beta/\phi = (1 - \chi)\eta_{net}$ gives $q_L = N_0/D_0$. In the ramp regime the coefficient becomes $\tilde{\sigma} + k$ and there is an extra term $-kq_F$, so $(D_0 - k)q = N_0 - kq_F$ and $q_S = (kq_F - N_0)/(k - D_0)$. Along (21), $F(q) \equiv q^d(q) - q = (k - D_0)(q - q_S)$ in the ramp regime, so the ceiling state $q_H = q_B$ is admissible exactly when $q_B \geq q_S$. With the calibration, $N_0 = 0.5244$, $q_L = 1.20$, $q_S = (1.05 \times 2.2 - 0.5244)/0.613 = 2.913$.

B. Robustness: The Wage-Led Regime

All parameters are those of Table 3 except η_{net} . Two versions are shown, because the effect of η_{net} depends on what is held fixed: (A) keeps autonomous demand g_A^0 at its benchmark, so q_L varies; (B) re-sets g_A^0 so that $q_L = 1.2$ in every row.

	η_{net}	Λ_{net}	D_0	ς needed	(A) fixed demand		(B) fixed $q_L = 1.2$	
					q_L	q_S	q_S	high state?
0.18 (benchmark)	0.38	0.133	0.367	> 0.52	0.82	2.94	2.74	yes
	0.063	0	0.437	> 0.62	1.20	2.91	2.91	yes
	0.00	0	0.500	> 0.71	1.45	2.88	3.11	yes
	-0.30	-0.105	0.605	> 0.86	1.76	2.80	3.56	barely
	-0.54	-0.189	0.689	> 0.98	1.93	2.71	4.11	no
	-0.86	-0.301	0.801	> 1.14	2.11	2.49	5.42	no

Table 7: The trap across regimes ($\varsigma = 1.5$, $q_F = 2.2$, $q_B = 3.6$). Positive η_{net} : profit-led; negative: wage-led.

The trap exists in the wage-led regime too, but the accelerator has to be stronger (the ς needed rises from 0.52 to 1.14); with the same autonomous demand (A) the low state sits higher and the trap is shallower, disappearing only in a strongly wage-led economy ($\eta_{net} < -1.07$ here); holding the low state fixed (B) the barrier is higher and for $\eta_{net} < -0.32$ it passes the ceiling, so there is no high state.

Running the slow dynamics for a wage-led economy ($\eta_{net} = -0.30$) next to the profit-led benchmark, with the push a three-year rise of 2 points in public demand growth:

	Profit-led $\eta_{net} = 0.18$	Wage-led (A) same demand	Wage-led (B) same $q_L = 1.2$
Low state q_L	1.20	1.76	1.20
Barrier $q_S - q_L$	1.71	1.05	2.36
Years in the trap before the high state is lost	about 20	about 23	under 1
Latest start year at which the push works	13	21	none
Smallest push that works at year 5 (pp)	1.5	0.8	none
Guarantee that ends the trap, s^*	1.00	0.45	1.00
Public demand that ends it, g_G^* (pp)	0.62	0.38	0.86

Table 8: Timing and policy in the two regimes.

In every case the wage share keeps rising in the trap, so waiting makes escape harder.

References

- Azariadis, Costas, and John Stachurski. 2005. "Poverty Traps." In *Handbook of Economic Growth*, Vol. 1, edited by Philippe Aghion and Steven N. Durlauf, 295–384. Amsterdam: Elsevier.
- Bhaduri, Amit, and Stephen Marglin. 1990. "Unemployment and the Real Wage: The Economic Basis for Contesting Political Ideologies." *Cambridge Journal of Economics* 14 (4): 375–393.
- Central Statistics Office (Ireland). 2025. "Productivity in Ireland Quarter 2 2025." Dublin: CSO. Released October 23, 2025.
- Central Statistics Office. 2026a. "Annual National Accounts 2025 and Quarterly National Accounts and International Accounts Q1 2026." Press statement. Dublin: CSO.
- Central Statistics Office. 2026b. "International Accounts Q4 2025: Balance of International Payments." Dublin: CSO.
- Cesaratto, Sergio. 2013. "The Implications of TARGET2 in the European Balance of Payments Crisis and Beyond." *European Journal of Economics and Economic Policies: Intervention* 10 (3): 359–382.
- Dixit, Avinash K., and Robert S. Pindyck. 1994. *Investment under Uncertainty*. Princeton, NJ: Princeton University Press.
- Economic and Social Research Institute. 2025. "Hare or Tortoise? Productivity and Growth of Irish Domestic and Foreign-Owned Sectors." *Economic and Social Review*.
- Freitas, Fabio, and Franklin Serrano. 2015. "Growth Rate and Level Effects, the Stability of the Adjustment of Capacity to Demand and the Sraffian Supermultiplier." *Review of Political Economy* 27 (3): 258–281.
- Goodwin, Richard M. 1951. "The Nonlinear Accelerator and the Persistence of Business Cycles." *Econometrica* 19 (1): 1–17.
- Goodwin, Richard M. 1967. "A Growth Cycle." In *Socialism, Capitalism and Economic Growth*, edited by Charles H. Feinstein, 54–58. Cambridge: Cambridge University Press.
- Hansen, Bruce E. 2000. "Sample Splitting and Threshold Estimation." *Econometrica* 68 (3): 575–603.
- Harrod, Roy F. 1933. *International Economics*. London: James Nisbet and Cambridge University Press.

- Harrod, Roy F. 1939. “An Essay in Dynamic Theory.” *Economic Journal* 49 (193): 14–33.
- Hicks, John R. 1950. *A Contribution to the Theory of the Trade Cycle*. Oxford: Clarendon Press.
- International Monetary Fund. 2023. “Boom Without Disease? Impact of Multinational Enterprises in Ireland.” IMF Country Report No. 23/412, Selected Issues. Washington, DC: IMF.
- International Monetary Fund. 2026. “IMF Executive Board Concludes 2026 Article IV Consultation with Ireland.” Press Release No. 26/225, 29 June.
- Timoney, Kevin. 2023. “Demystifying Ireland’s National Income: A Bottom-Up Analysis of GNI* and Productivity.” Dublin: Irish Fiscal Advisory Council Working Paper.
- Irish Fiscal Advisory Council. 2026. *Fiscal Assessment Report, June 2026*. Dublin: IFAC.
- Kaldor, Nicholas. 1966. *Causes of the Slow Rate of Economic Growth of the United Kingdom*. Cambridge: Cambridge University Press.
- Kalecki, Michal. 1971. *Selected Essays on the Dynamics of the Capitalist Economy, 1933–1970*. Cambridge: Cambridge University Press.
- Kaminsky, Graciela L., Carmen M. Reinhart, and Carlos A. Végh. 2004. “When It Rains, It Pours: Procyclical Capital Flows and Macroeconomic Policies.” *NBER Macroeconomics Annual* 19: 11–53.
- Kahn, Richard F. 1931. “The Relation of Home Investment to Unemployment.” *Economic Journal* 41 (162): 173–198.
- Keynes, John Maynard. 1936. *The General Theory of Employment, Interest and Money*. London: Macmillan.
- Nelson, Richard R. 1956. “A Theory of the Low-Level Equilibrium Trap in Underdeveloped Economies.” *American Economic Review* 46 (5): 894–908.
- Organisation for Economic Co-operation and Development. 2026. “Labour Productivity in OECD Economies: Recent Trends and Shifting Drivers of Growth.” In *OECD Compendium of Productivity Indicators 2026*. Paris: OECD Publishing.
- Serrano, Franklin. 1995. “Long Period Effective Demand and the Sraffian Supermultiplier.” *Contributions to Political Economy* 14: 67–90.

- Sinn, Hans-Werner, and Timo Wollmershäuser. 2012. “Target Loans, Current Account Balances and Capital Flows: The ECB’s Rescue Facility.” *International Tax and Public Finance* 19 (4): 468–508.
- Steindl, Josef. 1952. *Maturity and Stagnation in American Capitalism*. Oxford: Basil Blackwell.
- Thirlwall, Anthony P. 1979. “The Balance of Payments Constraint as an Explanation of International Growth Rate Differences.” *Banca Nazionale del Lavoro Quarterly Review* 32 (128): 45–53.
- Thirlwall, Anthony P. 2011. “Balance of Payments Constrained Growth Models: History and Overview.” *PSL Quarterly Review* 64 (259): 307–351.
- Thirlwall, Anthony P., and M. Nureldin Hussain. 1982. “The Balance of Payments Constraint, Capital Flows and Growth Rate Differences between Developing Countries.” *Oxford Economic Papers* 34 (3): 498–510.
- Tornell, Aaron, and Philip R. Lane. 1999. “The Voracity Effect.” *American Economic Review* 89 (1): 22–46.
- Verdoorn, Petrus J. 1949. “Fattori che Regolano lo Sviluppo della Produttività del Lavoro.” *L’Industria* 1: 3–10.