THE CITIZENS STANDARD

A Macroeconomic Model of a Two-Circuit Monetary System

Price Determination, Asset-Market Dynamics, and Stabilization under Citizen-Anchored Issuance

Neo-Solon
Neo-Solon@hotmail.com

WORKING PAPER — JUNE 2026

COMPANION PAPERS

The Citizens Standard: One Model, Many Systems (Neo-Solon, 2026a) · SSRN 6702518
The Citizens Standard: A Historical Counterfactual (Neo-Solon, 2026b) · SSRN 6735078
The Citizens Standard: Transition Architecture and Migration Mechanics (Neo-Solon, 2026c) · SSRN 6810741
The Citizens Standard: A Statutory Implementation Pathway (Neo-Solon, 2026d) · SSRN 6873798
The Citizens Standard: Full-Reserve Banking and the Two-Circuit System (Neo-Solon, 2026f) · SSRN 6939498
The Citizens Standard: External Interoperability and the Common Anchor (Neo-Solon, 2026g) · SSRN 6939600
The Citizens Standard: The Structural Buyer (Neo-Solon, 2026h) · SSRN 6945320
The Citizens Standard: The Issuance Engine (Neo-Solon, 2026i) · SSRN 6973261
The Citizens Standard: Empirical Validation of the Transactional Aggregate (Neo-Solon, 2026j) · SSRN 6973298
The Citizens Standard: Crisis Behaviour and Failure Modes (Neo-Solon, 2026l) · SSRN 6973358
The Citizens Standard: Comparison with Alternative Monetary and Distributional Systems (Neo-Solon, 2026m) · SSRN 6973338
The Citizens Standard: Distribution and Wealth Inequality (Neo-Solon, 2026n) · SSRN 6973342

Abstract

This paper develops a macroeconomic model of the Citizens Standard, the constitutional monetary architecture set out in the preceding papers of this series. The model decomposes the money stock by *use* — a transactional circuit that prices goods and an asset circuit that holds claims on the productive economy — and derives the framework's macroeconomic behavior from that single structural device. Three results follow. First, K1/K2 issuance lands in the asset circuit as a portfolio-choice equilibrium rather than by assumption: a wealth-neutral forced sale has a zero direct propensity to consume, so newly created money raises asset values held by citizens rather than consumer prices, with a bounded second-order leak that shrinks as the universal equity floors mature (Proposition 1). Second, although the framework makes the issuing authority a large, price-insensitive purchaser of equity, forward returns settle at a bounded valuation premium rather than collapsing — the purchaser is a net buyer while the floors accumulate and a bounded net seller as they mature, with a deceased citizen's residual passing to an heir's floor rather than being sold — and firm-side equity supply responds to valuation (Proposition 2). Third, the architecture's distributional instrument — a direct, equal citizen dividend — is also its fastest stabilization instrument: under price-level-path targeting it operates as an automatic countercyclical stabilizer that reaches households without depending on the credit channel, while the slow ownership channels are correctly relieved of stabilization duty. The model is explicitly short-run non-neutral and makes contact with the standard New Keynesian framework rather than departing from it; it differs in the locus of money creation and the ownership of newly issued claims, not in the rigidity assumptions. A recurring structural feature is that the framework's strained window is its youth: the price-level leak, the household-saving channel, and the return-compression premium are all largest at launch and resolve over a single maturation horizon. Beyond these results the model adds a consumer-dividend channel (K3) that pays a growth-indexed, equal citizen dividend out of the same real-growth budget that funds the floors — bounded by the price-stability locus rather than stacked on top of it — and that, by diverting issuance away from index purchases, also caps the Proposition 2 premium directly. The model's load-bearing simplification — a partition of money by use — is an approximation rather than an institutional wall, and the paper is explicit about where that partition is porous and what is required to calibrate the three quantities the results leave open.

JEL classification: E21, E42, E51, E52, E58, G12
Keywords: Citizens Standard; two-pool monetary system; asset-market dynamics; distributional macroeconomics; automatic stabilizers; structural buyer

Contents

1. Introduction.......5
2. Relation to the standard framework.......5
3. The two-pool money architecture (core device).......6
    3.1 The decomposition and the stability condition.......6
    3.2 Microfoundation: why K1/K2 issuance lands in the asset circuit.......6
    3.2a Illustrative magnitude of the leak.......7
    3.3 Proposition 1 (two-pool price stability).......9
    3.4 The transition handoff (third, transitory pool).......9
    3.5 Throttle speed: why KI is a control instrument, not just a transfer.......9
    3.6 Local stability and the maturing-circuit condition.......9
    3.7 Forward-looking determinacy and the nominal anchor.......11
    3.8 Bank credit and the robustness of circuit separation.......12
4. The issuance rule (FDCA).......13
    4.1 K1 — the citizenship channel (all Modes).......13
    4.2 K2 — the growth channel (all Modes; calibration is Mode-specific).......13
    4.3 K3 — the consumer-dividend channel (all Modes; Mode-parameterized).......14
    4.4 KI — the inflation-gap channel (Mode C permanent; Mode Λ conditional).......15
    4.5 KT — the transition channel (Mode T only, transitory).......16
    4.6 The throttle, in the engine’s terms.......16
    4.7 The conditional transition damper (KI_T, Mode T).......17
5. Household block and distributional dynamics.......18
    5.1 Two-account households.......18
    5.2 Result 2 (the no-net-saving condition).......18
    5.3 When Result 2 fails — and why it is the same window as Section 3.2.......19
    5.4 Distributional reading.......19
    5.5 Labor supply and the wealth-effect objection.......19
    5.6 The welfare-optimal dividend share.......20
6. Asset-market dynamics under a permanent structural buyer (Proposition 2).......21
    6.1 The three forces.......21
    6.2 Proposition 2 (bounded compression).......22
    6.3 The honest concession: the accumulation phase.......23
    6.4 Consistency with the series baseline.......24
    6.5 Illustrative magnitude of the premium.......24
    6.6 The maturity question: a bounded rebalancing, not a sell-off.......25
    6.7 The realizable floor: a general-equilibrium account.......25
7. Crisis-mode operations and automatic stabilizers.......29
    7.1 The stabilization assignment.......29
    7.2 Demand collapse — the automatic stabilizer engages.......30
    7.3 Inflation surge — a real asymmetry.......30
    7.4 Supply shocks — the same dilemma everyone faces.......30
    7.5 Comparison to a conventional reaction function.......31
    7.6 A linearized dynamic system and impulse responses.......31

8. Open problems....32
Technical Appendix....34
    A.1 Notation....34
    A.2 Propositions (formal statements)....35
    A.3 Proof sketch — the rebalancing channel (A2)....37
    A.4 Proof sketch — the no-net-saving condition (A3 / Result 2)....38
    A.5 Calibration and falsification....38
    A.6 Demographic equity-flow model (dating Proposition 2)....39
    A.7 Proof sketch — labor supply and growth robustness....40
    A.8 Proof sketch — delayed feedback and the two-speed design....40
    A.9 Linearized dynamics and impulse responses (Proposition 6)....41
    A.10 Forward-looking determinacy (Proposition 7)....42
    A.11 The welfare-optimal dividend share (Proposition 8)....43
    A.12 Robustness of circuit separation to bank credit (Proposition 9)....44
References....46

1. Introduction

The four preceding papers in this series establish the Citizens Standard's institutional architecture, its empirical performance against US historical data, its transition path, and its statutory implementation. This paper supplies the macroeconomic model that underlies them: a system describing how output, prices, and the monetary aggregates co-move under the framework, and how the architecture responds to shocks.

The model is built from a single structural device — a decomposition of the money stock by *use* rather than by instrument — and the framework's principal macroeconomic properties are derived from it rather than assumed. Three of these properties are non-obvious and are the paper's main results: that newly created money raises citizen-owned asset values rather than consumer prices (Proposition 1); that making the issuing authority a permanent equity buyer produces a bounded valuation premium rather than a return collapse (Proposition 2); and that the framework's distributional instrument is simultaneously its fastest stabilization instrument (Section 7). Each falls out of the accounting once the two-circuit structure is written down carefully.

What this paper contributes. First, a formal two-circuit decomposition of the money stock by *use* rather than by instrument, and the steady-state price-stability condition that follows from it (Section 3.1). Second, a microfoundation showing that K1/K2 issuance lands in the asset circuit as a portfolio-choice result rather than an assumption, with a derived and bounded leak into consumer prices that shrinks as the universal floors mature (Section 3.2). Third, a formal price-stability proposition and its transition-period counterpart (Sections 3.3–3.4). Fourth, a result on stabilization speed showing that the K1/K2 issuance throttle is a slow stabilizer and the KI dividend a fast one, which reframes KI as a control instrument and not merely a distributional choice (Section 3.5), together with a convergence proposition giving the explicit condition under which the path-targeting dividend rule closes the price-path gap monotonically rather than oscillating (Proposition 3), and a conditional transition damper (`KI_T`, Section 4.7) that closes the residual credit contraction the transition paper's lending facility concedes it cannot, as additive issuance whose stability is inherited from that proposition. Fifth, a household block that derives the no-net-saving condition underlying price stability rather than assuming it (Section 5). Sixth, a formal treatment of asset-market dynamics under a permanent structural buyer, showing that return compression is a bounded valuation premium rather than a secular return collapse, because the structural buyer is also a structural seller in demographic steady state and firm-side equity supply responds through the Q-channel (Section 6, Proposition 2). Seventh, a consumer-dividend channel K3 (Section 4.3) that pays a growth-indexed universal dividend from the same budget that funds the floors — bounded by the price-stability locus rather than added on top of it — and that, by diverting issuance away from index purchases, doubles as a direct throttle on the Proposition 2 premium. Section 7 then applies the stabilization-speed result to crisis operations; a recurring theme across the propositions is that the framework's strained window is its youth, with price-leak, worker-saving, and return-compression channels all resolving on a single maturation horizon.

2. Relation to the standard framework

The model is deliberately written to make contact with the New Keynesian consensus rather than to stand apart from it. Money is non-neutral in the short run here, through the same nominal-rigidity channel assumed in standard models; the framework does not rest on short-run neutrality and does not need to. Where it departs is in the *locus* of money creation and the *ownership* of newly issued claims, which is what generates the two-circuit structure. The quantity-theoretic relation used below for the transactional circuit is a modeling convenience for the steady-state results, not a monetarist commitment to a stable velocity at all horizons; the short-run dynamics are governed by the stabilization analysis of Section 3.5 and the crisis operations of Section 7.

3. The two-pool money architecture (core device)

3.1 The decomposition and the stability condition

Total money M_t is decomposed **by use**, not by instrument:

- M^T_t — the *transactional* circuit: money used for wages, goods, and commerce.
- M^A_t — the *asset* circuit: money used to hold and exchange claims on the productive economy.

with $M_t = M^T_t + M^A_t$.

The price level for goods is governed by the transactional circuit only:

$$P_t = M^T_t \cdot V / Y_t$$

with V the (stable) transactional velocity and Y_t real output. The central architectural claim is that K1/K2 issuance is directed into M^A (the asset circuit) and is held there by the portfolio behavior of asset-holders, so that goods-price stability requires M^T_t to grow with Y_t — which is what the issuance rule and the liquidation flow jointly deliver in steady state. The formal statement is Proposition 1.

3.2 Microfoundation: why K1/K2 issuance lands in the asset circuit

The old assumption (A2) — “proceeds stay in the asset circuit” — should not be an axiom; it can be derived from the seller’s portfolio choice.

Consider the agents who sell shares to the FDCA: unconstrained wealth-holders in the asset circuit. Each holds money m and risky assets e, with wealth $W = m + e$, and chooses a target money share $\mu^* = m/W$ solving the usual liquidity-preference trade-off, $\mu^* = \mu(r, \eta)$ — decreasing in the equity return r, increasing in money’s convenience yield η. Consumption follows a wealth rule, $c = \kappa_c \cdot W$ (the retiree drawdown rate, or the permanent-income fraction), so it depends on the *level* of wealth, not its composition.

The FDCA’s purchase forces a sale of Δ: money rises to $m + \Delta$, equity falls to $e - \Delta$, and W is unchanged. Two things follow mechanically:

- **Consumption:** $\Delta c = \kappa_c \cdot \Delta W = 0$. The proceeds are not spent on goods, because the swap left total wealth unchanged. This is (A2), now *derived*: the direct MPC out of a wealth-neutral asset swap is zero.
- **Portfolio:** the post-sale money share $(m + \Delta)/W$ exceeds μ^*, so the agent rebalances — buying Δ of assets to restore the target. The cash is redeployed inside the asset circuit.

Aggregating the rebalancing (a hot-potato across asset-holders), the Δ cannot leave the asset circuit: M^A rises by Δ, and for the aggregate to hold its target share μ^* of a now-larger asset stock, **asset prices must rise** until asset wealth grows by Δ/μ^*. The injection inflates asset values by a multiple $1/\mu^*$ of the money created and, to first order, leaves goods prices untouched. This is the price-stability result and the asset-appreciation result at once — and the appreciation accrues to every equity holder, including the locked Stable Floors.

The honest second-order term. Rising asset prices are a capital gain to whoever holds the assets, and unconstrained holders can spend out of that gain (the standard asset-wealth effect), so a fraction leaks into goods demand:

$$\text{Leak}_t = \kappa_W \cdot (\Delta_t / \mu^*) \cdot (1 - s_t)$$

where κ_W is the MPC out of asset wealth (empirically ≈ 0.02–0.05) and s_t is the share of risky assets held inside *locked* floors, whose holders cannot spend the gain. This is the only residual through which K1/K2 issuance touches the price level — small, bounded, and **shrinking over time**: as universal floors come to hold a larger share of total equity ($s_t \to 1$), more of the appreciation is trapped in accounts that cannot spend it, and the leak vanishes. *The framework becomes more price-stable as it matures.* So (A2) is not an axiom but the mature limit ($s_t \to 1$) of a derived relationship, with a quantified deviation for the immature system.

The endogeneity of μ^* (a second-order term that cuts in the framework's favor). The leak formula above holds the target money share μ^* fixed, but μ^* is itself a function of the equity return, $\mu^* = \mu(r, \eta)$, decreasing in r. The same appreciation that drives the leak — asset prices rising by Δ/μ^* — lowers the forward return r, which *raises* the desired money share μ^*. A higher μ^* shrinks the appreciation multiplier $1/\mu^*$, and therefore shrinks the Δ/μ^* base on which the leak is computed. Writing the return sensitivity as $\mu_r \equiv -\partial\mu^*/\partial r > 0$ and the appreciation-to-return passthrough as $r_Q < 0$, the first-order correction to the money share is $\Delta\mu^* \approx -\mu_r \cdot r_Q \cdot (\Delta/\mu^*) / W_A > 0$ (a rise), so the corrected leak is

$$\text{Leak}_t = \kappa_W \cdot (\Delta_t / (\mu^* + \Delta\mu^*)) \cdot (1 - s_t), \text{ with } \Delta\mu^* > 0,$$

which is strictly smaller than the fixed-μ^* value. The effect is second-order and we do not lean on it for the headline magnitude — the illustrative bound of §3.2a deliberately uses the *un*corrected, larger fixed-μ^* leak as the conservative case. The point is only that endogenizing μ^* moves the residual the right way: the framework's own price effect induces a portfolio shift toward money that partially absorbs that same effect. μ_r is the one new quantity this introduces; it is bracketed, not pinned, and enters the falsification set (A.5) as the sign restriction $\mu_r > 0$ rather than as a point value — a wrong-signed μ_r (money demand rising with returns) would remove this favorable correction but, because the headline bound already omits it, would not loosen the §3.2a result.

> **Calibration note (κ_W).** The 0.02–0.05 range is the mainstream estimate for the MPC out of stock-market wealth. Poterba (2000) calibrates the wealth effect at roughly one to two cents per dollar in the year after a change in stock values; Chodorow-Reich, Nenov & Simsek (2021), using a local-labor-market design, estimate about 3.2 cents per dollar per year; Brayton & Tinsley (1996) use ≈ 0.030 out of equity (vs. ≈ 0.075 out of other net worth) in the FRB/US model. The estimate is contested — Dynan & Maki (2001) find values as high as 0.05–0.15 for households with moderate securities holdings, with the most likely gain in the lower part of that range — and, crucially, it is highly heterogeneous across the wealth distribution: Di Maggio, Kermani & Majlesi (2020) estimate an MPC out of unrealized capital gains of roughly 0.23 for the bottom 50 percent but about 0.03 for the top 30 percent. This heterogeneity supports the use of a low κ_W here rather than undermining it: the agents executing the rebalancing in this section are the unconstrained, asset-rich sellers — precisely the top-of-distribution households whose MPC out of asset wealth is lowest. High-MPC households hold little equity and so contribute little to the Δ/μ^* appreciation base. A low κ_W is therefore the empirically appropriate value for the population that actually holds the relevant wealth, not a convenient assumption.

This separation also resolves two objections that are usually conflated. K1/K2 issuance shows up as **asset-price** inflation — which citizens own, through the floor — not as **consumer-price** inflation, which is the small and shrinking leak. The framework does not deny that issuance has a price effect; it locates that effect in asset values held by citizens rather than in the cost of living.

3.2a Illustrative magnitude of the leak

The leak formula is bounded in principle; this subsection puts an order of magnitude on it using only figures already on record, so a reader can see how large the consumer-price effect could plausibly be. This is a *worked illustration, not a calibration*: the issuance figures and the κ_W range are anchored, but μ^* and the s_t path are assumed within plausible bands and are exactly the quantities a full empirical calibration (deferred to the empirical paper) would pin down. A reader who prefers different assumptions can recompute directly from the inputs below.

Table 1. Price-leak calibration inputs (Section 3.2).

Input	Symbol	Value used	Status
M2 money supply	M2	\$22,366B	Anchored (Neo-Solon 2026a, A.1)
Mode B annual issuance	Δ	≈ \$447B	Anchored (Neo-Solon 2026a, §5.1)
Mode A annual issuance	Δ	≈ \$78B	Anchored (Neo-Solon 2026a, Table 2)
MPC of asset-rich sellers	κ_W	0.03 (band 0.02–0.05)	Literature-anchored (§3.2 note)
Target money share	μ^*	0.15 (band 0.10–0.25)	Assumed
Locked-float share at launch	s_t	0 (worst case)	Structural (no matured floors at launch)

Evaluating Leak = $\kappa_W \cdot (\Delta/\mu^*) \cdot (1 - s_t)$ at launch ($s_t = 0$, the most adverse point). Under Mode B’s 60/40 split the asset-circuit injection Δ is the FDCA equity purchase of \$272 billion — the 40 percent paid as the K3 dividend enters the transactional circuit directly (the Section 4.3 consumer-dividend channel) and is therefore not an asset-circuit leak:

- **Mode B (price-stable), central case** ($\kappa_W = 0.03$, $\mu^* = 0.15$): leak ≈ \$54.4B ≈ **0.24% of M2**.
- **Mode B, full assumption band**: leak ranges ≈ \$22B–\$136B, i.e. **0.10%–0.61% of M2**, with the high end requiring both a high MPC (0.05) and an implausibly low money share (0.10).
- **Mode A (deflationary), central case**: leak ≈ \$16B ≈ **0.07% of M2**, an order of magnitude smaller because Mode A issues far less.

The leak then declines mechanically as floors mature and s_t rises, holding the central case fixed: from 0.24% of M2 at $s_t = 0$, to 0.18% at $s_t = 0.25$, 0.12% at $s_t = 0.50$, 0.06% at $s_t = 0.75$, and 0.02% at $s_t = 0.90$. The reading: even in the price-stable mode and even at launch with no maturity buffer, the consumer-price leak is bounded below roughly 0.6% of M2 across the full plausible parameter space, sits near 0.24% in the central case, and falls toward zero over the maturation horizon. What a full calibration would add is an empirically grounded μ^* and a dated s_t path; the order of magnitude does not depend on those refinements. The figure above is the asset-circuit spillover only (approximately \$54.4 billion, 0.24 percent of M2). Under Mode B’s 60/40 split the complete transactional-circuit picture adds the consumer dividend, which is paid directly into the goods circuit: the K3 dividend is approximately \$175.3 billion (0.78 percent of M2), so the combined transactional injection is approximately \$229.7 billion, about 1.03 percent of M2 (\$54.4 billion + \$175.3 billion). This total is bounded by the price-stability locus rather than stacked on top of it. Because K3 is carved from the same $g_r \cdot M2$ budget (Section 4.3) and substitutes for the liquidation flow that the pure-floor configuration would release only as floors mature, M^T continues to grow with real output Y_t and goods-price drift remains approximately zero. The 60/40 split therefore brings the transactional circuit onto the price-stability locus at launch, where the pure-floor configuration reaches it gradually through the maturing drawdown.

3.3 Proposition 1 (two-pool price stability)

Assumptions. (A1) $M = M^T + M^A$, $P_t = M^T_t \cdot V / Y_t$, with V stable. (A2) the seller-rebalancing result of Section 3.2: K1/K2 issuance enters M^A and, to first order (the mature limit $s_t \to 1$), does not enter M^T. (A3) households in the constrained regime — no net worker saving out of the transactional circuit into the asset circuit. This is not stipulated: it is derived as Result 2 (Section 5.2) from a life-cycle worker whose equity-saving motive is already satisfied by the locked Stable Floor, and it holds in the mature system on the same horizon over which (A2) tightens.
Claim. Under (A1)–(A3), if the transactional circuit grows with real output, $M^T_t / M^T_{t-1} = Y_t / Y_{t-1}$, then the goods price level P_t is constant. K1/K2 issuance of any size consistent with the rule affects asset prices (via $1/\mu^*$) and the size of the Stable Floors, but not the cost of living, except through the bounded leak of Section 3.2.
Mechanism. The transactional circuit is fed by the liquidation flow L_t — the conversion of matured floor balances into spendable income — and by wages. The issuance rule and the liquidation flow are jointly calibrated so that M^T tracks Y. Issuance beyond that tracking requirement is absorbed by the asset circuit, where it raises claims values rather than goods prices.

3.4 The transition handoff (third, transitory pool)

During the migration there is a third, *transitory* pool: bank-created money under the legacy system, M^T_{bank}, which the Transition Lending Facility (TLF) progressively replaces. This makes the size of the TLF a *derived* quantity, not a free parameter: it must be large enough to offset the contraction of M^T_{bank} as banking separation proceeds, and no larger. As the transition completes, $M^T_{bank} \to 0$ and the system collapses to the two-pool steady state of Proposition 1. The transition is therefore not a separate regime with its own rules; it is the two-pool model with a vanishing third term.

3.5 Throttle speed: why KI is a control instrument, not just a transfer

What holds the economy on the price-stability locus? The throttle (Section 4) — but the levers differ sharply in speed. Today's liquidation flow L_t is set by floors accumulated *decades* earlier; it cannot be changed by adjusting issuance now. So the K1/K2 throttle is a **slow stabilizer**, acting only on future floor sizes with a decade-scale lag. High-frequency price control must run through a fast lever: the KI dividend (immediate) or active management of the liquidation rate. Pure **Mode B** (KI = 0) is price-stable in steady state for an economy at the one-half transaction-ratio balance point, but has weak short-run stabilization, relying on L_t being on-target by balanced-growth construction. This is a structural argument — straight from the two-pool accounting — for why the framework needs the KI channel and the adaptive Mode Λ as genuine *control* instruments, not merely distributional choices.

3.6 Local stability and the maturing-circuit condition

Proposition 3 established convergence of the path-targeting rule conditional on a reduced-form pass-through ψ, with the condition $\psi\lambda < 1$, but left ψ free. The two-pool accounting of Section 3.1 pins it down — converting Proposition 3 from a conditional convergence statement into a determinacy result with an explicit, observable, and time-varying stability ceiling, and fixing the money base against which the rule must be written.

Structural pass-through. Write the gap-closure injection against a money base B: it deposits $\lambda \cdot x_{t-1} \cdot B$ spendable dollars into the transactional circuit. Since $P_t = M^T_t \cdot V/Y_t$ (A1), an injection of that size raises M^T, and hence the price level, by the factor $\lambda \cdot x_{t-1} \cdot (B/M^T)$. The realized gap-closing inflation is therefore $\pi^{KI}_t = \lambda \cdot x_{t-1} \cdot \psi$ with $\psi = B/M^T$, and the gap law of Proposition 3 becomes $x_t = (1 - \psi\lambda)x_{t-1} + \varepsilon_t$. The pass-through is not a free parameter: it is fixed by the base against which the injection is scaled.

Stability and a scope note. The homogeneous gap dynamics have a single root $r = 1 - \psi\lambda$, so the path-targeting feedback is locally stable iff $|1 - \psi\lambda| < 1$, i.e. $0 < \psi\lambda < 2$, and monotone iff $\psi\lambda < 1$. This is the path-rule counterpart of the Taylor principle with the sign inverted: an interest-rate rule needs the response bounded below (more than one-for-one), whereas a full-pass-through quantity rule on a price path needs it bounded above. Because the gap law here is backward-looking, this is the stability of an adaptive feedback rather than a rational-expectations determinacy result. What would underwrite determinacy in a forward-looking setting is the rule's object: by targeting the cumulative price-level path (Section 4.4) rather than a one-period inflation rate, it pins the price level to a deterministic path and supplies a nominal anchor, where an interest-rate peg or an unindexed money-growth rule would leave the level free. A full forward-looking determinacy analysis is developed in Section 3.7 and Appendix A.10 (Proposition 7); what follows first establishes the stability and gain-selection properties of the feedback itself.

The base is a stability question. Indexing the gap-closure term to broad money, $B = M2$, gives $\psi = M2/M^T = 1 + M^A/M^T$. At the paper's illustrative near-boundary value $\psi = 1.8$ (the Appendix replication of Proposition 3), $\psi\lambda = 0.90$ — inside the monotone region but close to its edge; a narrower reading of the transactional circuit would raise ψ and shrink that margin further. Either way the framework's central thesis — that K1/K2 issuance accumulates in a large and growing asset circuit — means M^A/M^T rises over the life of the system, so under M2-indexing ψ rises above its launch value and the margin $\psi\lambda$ climbs toward the oscillation onset at $\psi\lambda = 1$ and, with continued maturation, beyond it (Figure 1). The instability is latent at launch and realized by maturity.

Resolution. Stability across the system's life therefore requires indexing the gap-closure injection to the circuit it actually inflates, $B = M^T$ — the form adopted in Section 4.4. This sets $\psi \equiv 1$, so $\psi\lambda = \lambda$ and the ceiling becomes the fixed, maturity-invariant $\lambda < 1$, with the baseline $\lambda = 0.5$ comfortably inside at every stage of the system's life. (Equivalently, an M2-based rule with an adaptive gain $\lambda_t = c \cdot M^T_t/M2_t$ holds $\psi\lambda = c$ constant; M^T-indexing is preferred because it requires no real-time money-ratio estimate and states the rule in the units of the quantity it controls.) The transactional-circuit base is not a stylistic choice but the condition under which the price-path rule stays stable as the asset circuit grows.

Gain selection: the variance floor. The stability range $0 < \psi\lambda < 2$ governs whether the gap converges, not how tightly it is held under ongoing shocks. With i.i.d. innovations ε_t of variance σ^2, the stationary gap variance is $\mathrm{Var}(x) = \sigma^2/[1 - (1 - \psi\lambda)^2]$, which is U-shaped in the gain: it diverges as $\psi\lambda \to 0^+$ (no correction) and as $\psi\lambda \to 2^-$ (near-unstable oscillation), and attains its floor σ^2 at $\psi\lambda = 1$, where the root is zero and the rule closes the expected gap in a single period. The variance-minimizing gain therefore lies on the monotone–oscillation boundary, which sharpens the case for transactional-circuit indexing. Under $B = M^T$ the gain is $\psi\lambda = \lambda$ with $\psi \equiv 1$ fixed and known, so the operator can raise λ toward one — approaching the variance floor — at no maturity risk; the baseline $\lambda = 0.5$ (variance $\approx 1.33\, \sigma^2$) is a deliberately conservative point on that locus, and tightening toward one is safe and monotone. Under $B = M2$ the same tightening is unavailable: ψ drifts upward with maturity, so a gain set near the variance floor at launch is carried through the oscillation boundary and toward divergence as the asset circuit grows. The near-boundary illustration $\psi = 1.8$, $\psi\lambda = 0.90$ sits within one percent of the floor precisely because it is already close to the unstable edge — a tight-but-fragile point that transactional indexing converts into a tight-and-robust one.

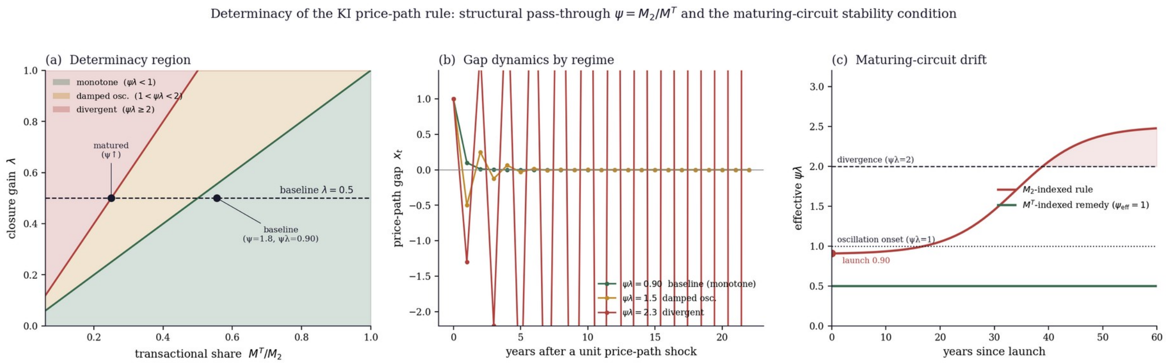

Figure 1. Determinacy of the KI price-path rule. (a) Stability region in the gain–money-share plane: monotone below λ = M^T/M2 (green), damped-oscillatory between M^T/M2 and 2M^T/M2 (gold), divergent above (red); the baseline λ = 0.5 is stable for M^T/M2 > 0.25 and monotone for M^T/M2 > 0.5. (b) Simulated price-path gap after a unit shock for the three regimes. (c) Effective ψλ over the system's life: an M2-indexed rule (red) drifts above the divergence threshold as the asset circuit matures, while the M^T-indexed rule adopted here (green) holds ψλ = λ = 0.5 at all maturities.

The two speeds, formalized. Section 3.5 distinguished the slow lever (the liquidation flow L_t, fixed by floors accumulated decades earlier) from the fast lever (KI). The determinacy analysis makes the division exact: L_t is predetermined at every date — it responds to floors set a generation ago, not to the current gap — so it enters the gap law as exogenous forcing, not feedback, and contributes no root to the characteristic equation. Determinacy is governed entirely by the contemporaneous feedback gain ψλ. This is why the decade-scale lag in L_t — destabilizing in any contemporaneous feedback loop, long delays being the classic source of feedback instability — is harmless here: the slow lever is a level, not a controller. Had the design instead made L_t respond to the current gap with its inherent lag, the system would carry a delay term and the stable region would collapse. The framework's separation — slow lever predetermined, fast lever in feedback — is the configuration that keeps the long lag out of the stability condition.

Why the slow lever is feedforward. The stability result of this section treats the decade-scale liquidation flow L_t as a predetermined schedule — forcing, not feedback. That choice is forced by the delay, not adopted for convenience. Were the slow lever instead to lean on the lagged gap, the recursion would acquire a delayed feedback term, x_t = (1−ψλ)x_(t−1) − λ_L x_(t−d), and its characteristic polynomial z^d − (1−ψλ)z^(d−1) + λ_L would carry d−1 extra poles. Two things follow as the delay d grows (Proposition 5; proof and Figure 3 in A.8): the gain that preserves stability collapses toward the delay-independent small-gain bound |1−ψλ| + |λ_L| < 1, and — the decisive point — even inside the stable region the damping of the slowest mode vanishes like |ln λ_L|/d, so a decade-scale lever leaves the loop technically stable but arbitrarily sluggish, ringing at the delay period for years after a shock. A slow instrument therefore cannot do useful feedback work; committing it as a predetermined schedule preserves the clean first-order dynamics on which the rest of this section rests and leaves the fast KI throttle as the sole feedback channel.

3.7 Forward-looking determinacy and the nominal anchor

The scope note, discharged. Section 3.6 established that the path-targeting feedback is locally stable as a backward-looking adjustment, and flagged the complementary forward-looking question — whether rational expectations pin a unique equilibrium price path, or whether self-fulfilling beliefs (sunspots) can arise — as open. This subsection settles it. Add a forward-looking money-demand block: real transactional balances fall with expected inflation, m^{T}_t − p_t = −α E_t[p_(t+1) − p_t] + (output term), the standard Cagan form. This is not an ad hoc identity: it is the first-order log-linear form that money-in-utility and shopping-time models deliver from household optimization (Sidrauski 1967; Lucas 2000; Ireland 2009), and α is the semi-elasticity of money demand with respect to expected inflation. The constitutional object here is not this demand equation but the KI quantity rule that operates against it; the demand side is inherited private behavior, not a rule we impose. With KI setting the transactional quantity in response to the price-path gap, the equilibrium gap obeys a single forward-looking equation E_t x_(t+1) = θ x_t + (shock) in the jump variable x_t.

Determinacy without a Taylor principle. The explosive root is $\theta = 1 + (1+\varphi)/\alpha$, where α is the money-demand semi-elasticity and φ the gap response. Because $(1+\varphi)/\alpha > 0$, $\theta > 1$ for every $\alpha > 0$ and every $\varphi \geq 0$: the price level is determinate even when the gap response is passive ($\varphi = 0$). The premise doing the work is $\alpha > 0$, which is among the most robustly established facts in monetary economics: the semi-elasticity is positive in every estimate from Cagan (1956) through the modern micro-founded literature, with US values below roughly 2 (Ireland 2009) and larger values only in high-inflation regimes. The sign of the result is therefore not in question. The margin is also comfortable at the framework's operating point: at $\alpha \approx 2$ with $\varphi = 0$, $\theta \approx 1.5$, half a unit clear of the determinacy boundary; even at a high-inflation $\alpha \approx 22$ the root is $\theta \approx 1.05$ — thin but still determinate. Determinacy is lost only in the limit $\alpha \to \infty$ ($\theta \to 1$), which no finite economy occupies. Because α rises as the nominal interest rate falls, the margin is thinnest in near-zero-nominal-rate environments; the framework's zero-inflation target paired with a normal positive real rate sits in the modern-US range rather than that limit, so the comfortable margin is the relevant case. This is the formal content of the 'inverted Taylor' remark of Section 3.6 — determinacy here does not require an aggressive reaction. It is supplied by the nominal anchor itself: because the Citizens Standard fixes the price level through the quantity of transactional money ($P\_t = M^T\_t \cdot V/Y$), a money-supply object, the level is pinned each period and rational expectations have no free dimension to coordinate on. An interest-rate instrument has no such anchor — the Fisherian root is φ, so determinacy demands the Taylor principle $\varphi > 1$ (Figure 5a). The two-circuit design buys determinacy that an interest-rate regime must earn with activism. Section 4.4 extends this anchor to a two-sided one: written symmetrically and drained by reverse-KT when the price level runs above the target path, KI pins the level from both sides, so the determinacy established here defends against above-target as well as below-target gaps.

Relation to the determinacy literature. The nearest antecedent is the banking New Keynesian model of Piazzesi, Rogers, and Schneider (2019), in which an endogenous convenience yield on inside money dampens interest-rate policy and delivers determinacy above a threshold even without the Taylor principle; the difference is one of channel and instrument. That model remains an interest-rate regime whose anchoring works through a liquidity premium, whereas the anchor here is the quantity of transactional money and the threshold governs asset-to-consumer spillover rather than money-supply rigidity. The result sits within a known family but is not reducible to its established members, and the distinction is worth stating precisely. Woodford (1995), in the paper whose title this section's most nearly echoes, shows the price level can be determinate without control of a monetary aggregate — but there the anchor is fiscal: the equilibrium level is the one that equates the real value of government liabilities to the present value of future surpluses (the fiscal theory of the price level). Here the anchor is monetary and contains no fiscal term: the determinacy root $\theta = 1 + (1+\varphi)/\alpha$ depends only on the transactional money-demand semi-elasticity α and the gap gain φ, with no government surplus, debt stock, or intertemporal budget condition entering. The two results are therefore complementary rather than competing, and can hold simultaneously. Equally, the finding is distinct from the price-level-targeting determinacy results (Giannoni 2014; Bauducco and Caputo 2020), in which an interest-rate rule attains determinacy below the Taylor bound by responding to the price level: those remain interest-rate regimes, whereas the anchor here is the quantity of transactional money itself. Finally, the result must answer the classic cautionary finding that constant-money-growth regimes can admit both perfect-foresight and sunspot indeterminacy (Woodford 1994, in the cash-in-advance setting). It escapes that case because KI is not a constant-growth rule but a quantity rule that responds to the price-path gap; it is the gap response, together with the segmentation that keeps asset-circuit money out of the transactional pricing relation, that removes the free dimension a constant-growth rule would leave open. What we believe distinctive here, and are not aware of a direct precedent for, is neither determinacy without an aggressive interest-rate reaction (long known) nor level anchoring as such, but the two-circuit route to it: a segmented money demand in which the transactional circuit pins the level while the asset circuit is walled off up to a quantified, design-controllable coupling threshold (Section 4.4). The separation this route relies on is not only assumed but has direct empirical support: in the companion validation study (Neo-Solon 2026i), a transaction-active money aggregate carries a significant consumer-price signal in high-inflation regimes ($R^2 \approx 0.19$, robust to a PCE index) while broad M2 does not, and in an encompassing regression the transactional aggregate stays significant as M2 drops out — the consumer-price side of the two-circuit split showing up directly in the data. The support is in-sample and regime-dependent rather than an out-of-sample forecasting edge, so it is offered as corroboration of the premise, not proof of it.

The asset circuit and the role of separation. With the asset circuit added as a second forward-looking variable — an asset-valuation gap obeying a present-value relation — the system has two jump variables, and Blanchard–Kahn determinacy requires two explosive eigenvalues. Both diagonal blocks are explosive (the price level by the money anchor, asset valuations by discounting and transversality), so for weak asset↔consumer coupling the count is met and the equilibrium is unique. Determinacy survives until the coupling reaches ≈ 0.13 (Figure 5c); at the leak implied by the consumer-price calibration (≈ 0.03) the system sits comfortably inside the determinate region. The circuit separation that delivers price stability (Proposition 1) and shock containment (Proposition 6) thus also delivers a unique rational-expectations equilibrium — three payoffs of one structural feature. Proposition 7 (Appendix A.2; proof and Figure 5 in A.10) states the result; α and the coupling are the empirical quantities (A.5).

Supplementary replication. The forward-looking determinacy analysis is developed in a companion two-asset New Keynesian DSGE, calibrated to textbook values (Galí 2015). Solved on the framework's actual nominal anchor — the quantity/price-path rule, not an interest-rate rule — it reproduces the determinacy-without-a-Taylor-principle result of this section: the explosive root is $\theta = 1 + (1+\varphi)/\alpha$, which exceeds one for every money-demand semi-elasticity $\alpha > 0$ and every gap response $\varphi \geq 0$, so the price level is determinate even under a passive gap response, because the money-supply anchor pins the level each period where an interest-rate rule would require $\varphi > 1$. On the price-stability side, an issuance-shock impulse is small under floor-weighting and large under a pure dividend (about 3.3 percent at the quantity-theoretic limit), with the asset-to-goods leak as the load-bearing parameter — the same boundary the credit-displacement module locates from the money-creation side, so three independent routes (the propositions here, that empirical module, and the DSGE) converge on it. The model's scope is precise: its rigorous content is the determinacy result, while the price response is the model's long-run quantity-theoretic property (validated against that limit) rather than a full short-run impulse response from a perturbation solution; and the stylized two-block system shows the direction of the coupling-determinacy trade-off, with the exact ~0.13 threshold supplied by the full two-block system of Appendix A.12. Runnable as the dsge-twocircuit module in the distribution replication package.

3.8 Bank credit and the robustness of circuit separation

The banking objection. The two-circuit architecture rests on a separation: transactional money M^T sets the price level, asset-circuit issuance lands in M^A, and the two are linked only by a small structural leak (Section 3.2a, ≈ 0.03). A natural objection is that bank credit could couple the circuits even under the framework's full-reserve rule: if banks lend against asset-circuit wealth and the proceeds are spent on goods, bank credit opens a channel from the asset circuit into M^T — re-mixing what the architecture separates. This subsection asks when that channel is large enough to matter.

A reduced-form credit channel. Let banks extend credit against pledgeable asset wealth at loan-to-value m, and let a share χ_c of new credit be spent into the transactional circuit. The induced asset→consumer coupling is then $\chi_c \cdot m \cdot \varphi_{liq}$, where φ_{liq} is the fraction of asset wealth that is liquid and can be pledged. This is the crux: the citizenship floors are locked — non-tradeable and, equally, non-pledgeable — so only the liquid fraction φ_{liq} backs credit. The credit channel inherits the very constraint that produces the demographic and labor results. The total asset→consumer coupling becomes $\lambda_{leak} + \chi_c \cdot m \cdot \varphi_{liq}$, and circuit separation — equivalently, the forward-looking determinacy of Proposition 7 — survives as long as this stays below the coupling threshold $\zeta^* \approx 0.13$.

The lock is decisive. Write the credit intensity as $\kappa_{bank} \equiv m \cdot \varphi_{liq}$. Under the conservative reading in which credit and asset values reinforce each other — a financial accelerator, so the coupling is symmetric — separation requires $\kappa_{bank} < 0.32$. The locked baseline is far inside: with most citizen wealth held in locked floors, the pledgeable fraction is small ($\varphi_{liq} \approx 0.15$), so breaking separation would require a loan-to-value above 2 — impossible. Were the floors instead fully pledgeable ($\varphi_{liq} = 1$), the same calculation breaks separation at a loan-to-value of just 0.32, an ordinary mortgage. The non-pledgeability of the floor is what keeps bank credit from re-mixing the circuits; it is the banking-channel counterpart of the lock that anchors the demographic and labor results (Figure 7, Appendix A.12).

Supplementary replication. A companion analysis bears on a related coupling question — not whether bank lending leaks the asset circuit into M^T (above), but whether the dividend issuance itself can be treated as net-neutral because it displaces credit banks would otherwise create (the double-claim objection). Running the issuance rule on the 1960–2025 series, the displacement a cash dividend would need to keep its price impulse under one percent is high and robust: about 89 percent in a pure-dividend configuration and 73 percent at a floor-weighted dividend share of 0.4, and it stays in the 70–95 percent range across defensible choices of the transactional ratio and tolerance. The empirical crowding-out literature finds displacement is partial and state-dependent (near zero in slack, as in 2008–09), well below that requirement, and the quantitative-easing episodes confirm that large injections are not displaced near one-for-one. Issuance neutrality therefore rests not on credit displacement but on the framework’s two structural features: the issuance is growth-matched (capped at realized growth) and most of it is routed to the asset-buying floor (M^a), so only the dividend share reaches the goods circuit, where the un-displaced portion of a large cash dividend would carry a bounded inflation risk. Floor-weighting is thus the structural mitigant, and the result supports a floor-weighted operating point and a modest dividend. Runnable as the credit-displacement module in the distribution replication package.

A second safeguard. The structural buyer reinforces this. Because asset demand under the Citizens Standard is dominated by the rule-bound K1/K2 buyer rather than by credit, the financial accelerator is damped: credit-fueled spending does not chase asset prices the way it does in a conventional collateral cycle. With the accelerator damped, the critical credit intensity rises to ≈ 1.7, so even a fully pledgeable, highly leveraged economy would retain separation. The lock and the structural buyer are independent protections, and the locked baseline sits comfortably inside both. A full intermediary-sector treatment — bank balance sheets, full-reserve credit supply, and the conversion contraction — is developed in the companion banking paper (Neo-Solon, 2026f); the claim here is only that endogenous credit does not, at any feasible leverage, defeat the separation. Proposition 9 (Appendix A.2; proof and Figure 7 in A.12) states the result.

4. The issuance rule (FDCA)

This section states the non-discretionary issuance rule formally, in the notation of Section 3, so that the throttle of Section 3.5 is fully specified. The rule is the one established in the architectural paper (Neo-Solon, 2026a, Sections 3–6 and Technical Appendix A.2), extended in this revision of the series by the consumer-dividend channel K3 (Section 4.3); the architectural paper carries K3 in step, so the two remain consistent. All channels distribute equally per citizen at the point of issuance; what distinguishes the regimes (Modes) is which channels are active and how they are calibrated.

Launch reference values. The calibrations below use the architectural paper's launch-year anchors: nominal GDP ≈ \$30,762B, population N ≈ 342M, GDP per capita ≈ \$90,000, M2 ≈ \$22,366B, base-case real growth $g_r \approx 2.0\%$/yr, population growth $g_p \approx 0.5\%$/yr, and accumulation horizon T = 65 years.

4.1 K1 — the citizenship channel (all Modes)

K1 is a one-time deposit triggered by each verified new citizen (birth or naturalization), calibrated as a fixed fraction α of GDP per capita:

$$K1(t) = \alpha \cdot GDP(t)/N(t), \text{ with } \alpha = 0.025.$$

At launch, K1 = 0.025 × \$90,000 ≈ \$2,250 per new citizen; aggregate $K1\_agg$ ≈ \$9B/yr (≈ 0.04% of M2). Naturalization at age a is pro-rated by remaining horizon: $K1\_nat(a) = K1 \cdot (65 - a)/65$. K1 routes through the capital-markets channel into the recipient's locked Stable Floor account — in the notation of Section 3, it enters the asset circuit M^A, and the cash paid to share-sellers is what reaches circulation.

4.2 K2 — the growth channel (all Modes; calibration is Mode-specific)

K2 is the annual growth-matched deposit, distributed equally to all living citizens' Stable Floor accounts. Its calibration is the single lever that distinguishes the price regime:

Mode A (mild deflation): $K2\_agg = \beta\_A \cdot g\_r \cdot M2 - K1\_agg$, with $\beta\_A = 0.175$ (capture of the full real-growth-matched rate, residual of K1). Launch: 0.175 × 0.02 × \$22,366B ≈ \$78B envelope, less K1 ⇒ K2_agg ≈ \$69B/yr (≈ \$203/citizen). Combined issuance is ≈ 0.35% of M2, but under the dual-circuit result of Section 3.2 only the bounded leak — ≈ \$16B, about 0.14% of the transactional circuit M^T — reaches goods prices; against ≈ 2% real growth on the M^T base ($P = M^T \cdot V/Y$), this leaves net structural deflation of ≈ 1.86%/yr.

Mode B (price stability): $K2\_agg = 1.0 \cdot g\_r \cdot M2 - K1\_agg$ (the full real-growth-matched rate, residual of K1). Launch: total = 1.0 × 0.02 × \$22,366B ≈ \$447B, so $K2\_agg$ ≈ \$438B/yr (≈ \$1,282/citizen). Combined K1+K2 ≈ 2.0% of M2, matched to ≈ 2% real growth ⇒ ≈ 0% price drift in growth years for an economy whose transaction-active money share μ sits near one-half (the fixed-split balance point; see the closed form below) (K1 continues in the rare contraction years, acting as a small automatic stabilizer).

The stabilizing split has a closed form. Setting derived CPI to zero on the inflationary side of the fixed 60/40 split and neglecting the small fixed-floor term k_1Agg, the price-stabilizing consumer-dividend share solves to

$$\boldsymbol{\kappa\_d^* \approx (\mu - 0.20) / 0.80,}$$

where μ is the transaction-active share of broad money (M^T / broad). The real-growth rate g cancels: it enters both the transactional injection and the deflationary pull that injection must offset, and the two scale together. Price stability under the fixed split is therefore a structural property of the monetary aggregate — the share of money held in transaction rather than time-deposit form — essentially independent of the growth rate. For μ near one-half the stabilizing share is near 0.40, which is why Mode B's fixed 60/40 split is near-stable for such economies; for a time-deposit-heavy economy with μ ≈ 0.28 it falls to κ_d* ≈ 0.10, and the fixed split at 0.40 instead produces positive drift. Holding CPI at zero across arbitrary μ requires solving κ_d (and, on the deflationary side μ > 0.51, the KI channel) to the target rather than fixing the split — the defining operation of Mode Ω (Architecture, Section 8). The balance point itself is $\mu^* = 0.52 - 0.32 \cdot k_1\mathrm{Agg}/(m2 \cdot g)$, which sits near 0.51 for the calibrated economies.

Mode C (modest inflation): K2 as in Mode A; the inflation regime is carried by KI below.

Like K1, K2 routes into the asset circuit *M^A* via the capital-markets channel.

Calibration note (Mode B full rate). This paper uses the full real-growth-matched Mode B calibration (1.0 · *g_r* · M2), consistent with the architectural paper's Section 5.1 and Table 3 and with the empirical paper (Neo-Solon, 2026b). The full rate is what delivers the *M^T*-tracks-*Y* condition of Proposition 1. *(The architectural paper's launch-parameter table, Technical Appendix A.1, states this full-rate calibration consistently with its A.2.2 and Section 5.1.)*

4.3 K3 — the consumer-dividend channel (all Modes; Mode-parameterized)

K3 is the first channel to enter the transactional circuit M^T directly, and the only permanent one (KI, the inflation-gap stabilizer of Section 4.4, is the other, but it is conditional). Like KI it deposits unlocked, immediately spendable dollars equally to every citizen; unlike KI — a path-targeting stabilizer that issues only to close an inflation gap (Section 4.4) — K3 is indexed to realized real growth, the same input K2 reads. It is the present-tense complement to K2: where K2 routes growth-matched issuance into locked Stable Floors (future wealth), K3 routes a share of the same growth-matched budget to citizens as current spendable income (a claim on present output). It fills the one gap the other channels leave open — a standing, universal, growth-indexed consumer dividend available in all Modes, as opposed to KI (inflation-gap-triggered, Mode C only) and the liquidation flow L_t (retirees only).

The issuance leash. The defining property is that K3 shares one budget with K2 rather than adding a second. The growth-matched issuance that holds the money stock on the price-stability locus is fixed at the full real-growth-matched rate; K1, K2, and K3 partition it:

K1_agg + K2_agg + K3_agg = g_r · M2, with K3_agg = κ_d · (g_r · M2 − K1_agg) and K2_agg = (1 − κ_d) · (g_r · M2 − K1_agg).

Here κ_d ∈ [0, 1) is the consumer share and the per-citizen dividend is K3_agg / N. Because the three channels always sum to the one budget, K3 cannot create net-new issuance above the line: raising κ_d reallocates growth from the asset circuit to the transactional circuit, it does not expand M. This is the same 100%-of-growth constraint that already pins K2 in Mode B (Section 4.2), now enforced jointly on the correlated pair (K2, K3) rather than on K2 alone.

The price leash. K3 carries a second constraint that K2 does not, because it lands in M^T. It is subject to the Proposition 1 price budget alongside the other spendable flows — the liquidation flow L_t and, where active, KI:

L_t + KI + K3 ≤ ΔM^T |(Y-growth),

the transactional-circuit expansion consistent with M^T tracking Y. Within this ceiling K3 is goods-price-neutral by the growth-matching argument of Section 3.2: the new dollars are citizens' claim on the incremental real output the economy actually produced, so they transfer a consumption claim rather than adding nominal demand against fixed supply. Pushing κ_d past the ceiling is not a free enlargement of the dividend — it is, by construction, a deliberate step into a modest-inflation regime, i.e. the mechanism by which a price-stable Mode is converted into a Mode-C-type Mode. Crossing the line is therefore always a Mode choice, never an accident.

K3 is therefore Mode-parameterized. Setting κ_d = 0 recovers each existing Mode unchanged (all growth-issuance to floors). A positive κ_d within the M^T ceiling reconfigures a Mode in place — Mode B, for instance, retains its approximate price stability (near the one-half transaction-ratio balance point) while paying a standing universal dividend — and a κ_d beyond the ceiling defines the transition to the next price regime. The parameter thus both generates new Modes and supplies the governed path between them. Like K2 and KI, K3 reads the smoothed Composite Productivity Index through the Adaptive Smoothing Rule; and κ_d is bounded above not only by the M^T ceiling but by retirement adequacy — diverting too much of the growth budget away from floors can pull the projected floor below the worker's target E* (Result 2), which sets the upper end of the usable range. Within that usable range the wealth-maximizing choice is κ_d = 0. At the central real equity return, over the accumulation horizon to age 65, the locked floor's forced, diversified compounding dominates a current dividend — by roughly 5.6× against a dividend that is consumed, and still ahead of one the citizen self-reinvests unless they replicate the floor's return net of frictions (a 1.5-point shortfall, for example, leaves the floor ahead by about 1.9×); the two coincide only under idealized frictionless reinvestment. A positive κ_d is therefore a liquidity tradeoff rather than an efficiency gain — wealth-neutral only in that idealized case and wealth-reducing to the extent the dividend is consumed or self-invested at a lower net return. The low-budget Modes accordingly set κ_d = 0, and a positive κ_d remains a ratifiable choice bounded above by the M^T ceiling and by retirement adequacy.

This is an individual, price-taker comparison: one citizen's locked floor against one citizen's dividend at a fixed market return. It does not aggregate. If every citizen retained the full flow at full budget (κ_d → 0 with K2 at the whole line), the floors would deepen the capital stock until the aggregate stake outran what the stock can hold and the realized return fell with it (§6.7) — so the universal, realizable wealth-maximizing κ_d is bounded strictly above zero by the feasibility condition. The bound is slack for the low-budget Modes, where the floor capture is small and κ_d = 0 sits safely inside it; it binds for the full-budget configuration, which is why the balanced Mode sets κ_d = 0.40 (a 60/40 floor–dividend split) rather than zero. Feasibility and the welfare argument of §5.6 thus push the same way: toward a positive dividend share.

Finally, K3 relaxes the accumulation-phase concession of Section 6.3. Every dollar paid out through K3 is a dollar of growth-issuance that does not purchase index shares, so gross FDCA absorption falls by κ_d · (g_r · M2 − K1_agg) and net absorption A_t with it. K3 is thus simultaneously the consumer dividend and a direct throttle on the structural-buyer premium of Proposition 2: the larger the citizen dividend, the smaller the valuation premium A*/φ and the shorter the accumulation phase. The channel that pays citizens current income is the same channel that defuses the "the buyer absorbs the whole market" critique.

4.4 KI — the inflation-gap channel (Mode C permanent; Mode Λ conditional)

KI is the inflation-gap channel into the transactional circuit *M^T*: like K3 (Section 4.3) it deposits unlocked, immediately spendable dollars equally to all citizens, but it is triggered by the price-level gap rather than by growth and — alone among the channels — is permitted to issue above the growth-matched line, which is how it produces inflation. It is calibrated by price-level *path* targeting, which closes the cumulative gap between actual CPI and a target path growing at *π** per year, rather than a single-year gap (the path-targeting design self-corrects rather than oscillating):

target_price_level(t) = (1 + π*)^t, with π* = 0.02   gap(t) =
ln[target_price_level(t−1) / actual_price_level(t−1)], with λ = 0.5 (gap → 0 on the
target path)   KI(t) = max(0, λ · gap(t) · M^T + [π* + g_r] · M^T − K3_agg − L_t)

The natural log makes the closure rate consistent across multi-year accumulated gaps. The gap-closure term is indexed to the transactional circuit M^T, not M2 — Section 3.6 shows this is required for determinacy as the asset circuit matures; the M2-indexed form stated in Paper 1, §6.1 (Neo-Solon, 2026) is the canonical simplification, identical at the steady state. At steady state the price-path gap is closed (gap → 0), so KI rests at its maintenance level $(\pi^* + g_r) \cdot M^T$ net of the K3 dividend and the floor spillover L — about 1.98% of M2, or 3.8% of M^T, ≈ $443B (≈ $108/citizen/month) at launch — and scales gently with the economy thereafter as a constant share of a growing M^T. On a clean launch the price level begins on its own target path, so the error-correction term is zero and KI issues at this maintenance level from year one; only if the economy opens below its target path does the gap-closure injection temporarily lift issuance above the maintenance level, receding to it as the path is recovered. KI is inactive (KI = 0) in Modes A and B.

The two-sided extension. As written, KI is one-directional: the max(0, ·) floor lets it raise a below-target price level but not lower an above-target one, leaving the slow K1/K2 throttle and real growth to undo overshoots. The post-transition steady state closes this asymmetry with the symmetric reverse-KT drain (transition paper, Neo-Solon 2026c). The same price-path gap that scales KI issuance when it is positive scales, when it is negative, a withdrawal of transactional money: the digital-currency authority sells short-term bills against the standing stock retained within the operational band after the transition and extinguishes the proceeds, contracting M^T rather than expanding it. Because the drain lands in the transactional circuit and never touches the locked floors, it moves the price-relevant aggregate without disturbing citizen capital. Dropping the max(0, ·) floor and writing the rule two-sided — the sign of the gap-closure term selecting injection versus drain — leaves the convergence condition of Section 3.6 ($\psi\lambda < 1$) and the forward-looking determinacy of Section 3.7 unchanged in form but strengthens their content. The explosive root $\theta = 1 + (1+\varphi)/\alpha$ that delivers Proposition 7 is untouched, but the nominal anchor it formalizes now pins the price level from both sides of the target path rather than from below alone: the level can no more drift persistently above the path than below it, and the linearization underlying the determinacy result holds globally rather than only on the side where KI is active. The two-circuit design thus supplies not merely a determinate price level but a symmetrically defended one — the contractionary half of the anchor supplied by reverse-KT at no cost to the floors. This determinacy is not a technical by-product but the result on which the framework's price claims rest: because it makes the targeted regime the unique equilibrium, it is what allows a constitution to select an inflation rate and actually obtain it rather than merely announce it, and it is what lets Mode C's issuance register as its intended two percent rather than as an indeterminate rate. Where those claims appear in the architecture and banking papers, they inherit their content from this proposition.

4.5 KT — the transition channel (Mode T only, transitory)

KT is not a citizen-distribution channel. During the migration it funds Legacy Trust retirement of pre-existing sovereign debt; it throttles to dormancy once the public debt stabilizes within its operational band (the transition paper sets this at approximately 30 to 60 percent of GDP), at which point Mode T lands automatically in the Mode B steady state (the two share the same price-stable calibration), with KT retained as a symmetric open-market instrument. In the notation of Section 3.4, KT drives the vanishing third pool $M^T_{bank} \to 0$; it is the policy instrument that drives that term to zero without contracting the transactional circuit. Detailed mechanics are in the transition paper (Neo-Solon, 2026c).

4.6 The throttle, in the engine’s terms

Sections 4.1–4.5 specify the levers referenced in Section 3.5. The mapping to stabilization speed is now explicit:

- **Slow levers — K1, K2.** Both feed locked Stable Floors via M^A. They affect the transactional circuit only later, through the liquidation flow L_t when those floors mature

— a decade-scale lag. They cannot move M^T this year, which is why they are slow stabilizers.

- **Standing fast lever — K3.** Like KI it enters M^T immediately and equally, so it shares KI's price-budget leash; unlike KI it is not gap-triggered but growth-indexed, drawing its share κ_d from the K2 budget (Section 4.3). It is a distributional lever first and a control instrument only at the margin — the dial that sets how much of each year's growth reaches citizens as current income rather than locked floors.
- **Fast lever — KI.** It enters M^T immediately and equally. This is the channel available for high-frequency price control, which is why Section 3.5 concludes that KI (and the conditional KI of Mode Λ) is a genuine control instrument and not merely a distributional choice.

K2, K3, and KI are all calibrated on the smoothed Composite Productivity Index via the Adaptive Smoothing Rule (default five-year rolling average, switching to two-year at turning points when the latest reading deviates by more than two percentage points until it returns within one), and Method A (minimum of smoothed GDP and CPI growth) biases K2 toward conservative issuance when the measures diverge. These are rules-based input-preparation steps, not discretionary judgments, and they bound how fast even the "fast" lever can be moved.

4.7 The conditional transition damper (KI_T, Mode T)

The transition runs three jobs. Two are specified in the transition paper (Neo-Solon, 2026c): KT retires legacy public debt (price-path calibrated, asset-swap, consumer-price neutral, throttling to dormancy once the public debt stabilizes within its ~30-60% operational band), and the Transition Lending Facility (TLF) performs targeted replacement lending for the most vulnerable credit channels as banking separation proceeds. The transition paper is explicit that the TLF is not a complete offset: its coverage of annual credit at risk is approximately 12–38% at full-rate K2, with the remainder left to non-bank-dependent corporate credit. The residual — the bank-dependent credit contraction the TLF does not reach — shrinks the transactional circuit exactly while the system is most fragile, pulling the realized price level below the target path.

Neither existing instrument is a price-path stabilizer for that residual: KT is pointed at debt and routes to bondholders, and the TLF is a scheduled level offset, not a feedback response to a realized gap. By the stabilization-speed result of Section 3.5, the only fast lever that can respond to a realized transactional-circuit contraction is KI. The third job is therefore a fast, conditional, automatic damper that fills the residual the TLF concedes it cannot, measured by deviation from the price path. That instrument is the conditional `KI_T`.

KI_T is the Section 4.4 path-targeting rule applied during Mode T, with three distinguishing properties. First, it is **additive net-new issuance**, never diverted from K1 or K2 — the transition paper's guarantee that citizen K1 and K2 flows are untouched throughout holds exactly; floors are never reduced to fund the damper. Second, it is **conditional with a dead band** δ, firing only against a genuine downward price-path gap rather than sub-threshold noise. Third, its **intensity** $\chi \in [0,1]$ is set at ratification. Extending the Section 4.4 rule (with π^* the ratified destination Mode's target — 0 into Mode B, 0.02 into Mode C):

$$\text{damped}(t) = \text{sign}(\text{gap\_closure\_T}) \cdot \max(0, |\text{gap\_closure\_T}(t)| - \delta) \cdot \lambda \; \text{KI\_T}(t) =$$
$$\max(0, \chi \cdot ([\text{damped}(t) + \pi^* + g_r] \cdot \text{M2} - K1_{agg} - K2_{agg}))$$

The three transition instruments then partition cleanly: KT retires debt (routing to bondholders), the TLF replaces vulnerable lending on a schedule (12–38%), and KI_T absorbs the residual contraction as a price-path feedback rule routed equally per citizen into circulation. In the language of conditional policy rules, `KI_T` is **in force throughout the transition but non-binding (slack) while the realized price level holds to target**, becoming **binding** only under a residual-contraction shortfall.

Its stability is inherited, not separately assumed. KI_T uses the same closure rule scaled by $\chi \le 1$, so its effective gain is $\chi\lambda$ and, by Proposition 3 (Appendix A.2), its stability condition $\psi \cdot \chi\lambda < 1$ holds on the same condition as the steady-state rule with strictly more margin; the dead band further suppresses sub-threshold response. The damper cannot oscillate on any setting the steady-state rule does not.

Because KI_T is additive, the frontier χ controls is not stability-versus-floors but transition price stability versus additional circulating issuance: a higher χ holds the price path more tightly at the cost of more net-new spendable money during transition, while a lower χ lets the economy absorb the residual as managed disinflation. Floors are identical under every χ. At $\chi = 0$ the damper is withdrawn and the TLF carries the offset alone; at $\chi = 1$ with small δ it opens on any genuine slip. Both χ and δ are Tier 2 ratifiable choices.

KI_T is self-extinguishing. Once banking separation completes (*M^T_bank* → 0, Section 3.4) and legacy debt is retired, the structural contraction impulse vanishes, the price-path gap ceases to run persistently negative from that source, and KI_T → 0 by its own max(0,·). The system lands in Mode T-stable — K1 and full-rate K2, with the split solved for zero derived inflation (this is Mode Ω; it coincides with the fixed Mode B split at the one-half transaction ratio and is the Ω solve otherwise), with KT and KI_T both withdrawn from the configuration. The damper is a transitory instrument that exits structurally as separation completes, not a permanent channel, unless the society subsequently ratifies a steady-state Mode (C or Δ) that retains a KI.

5. Household block and distributional dynamics

This section supplies the household optimization that Proposition 1 leans on. Assumption (A3) — that workers do not, on net, divert the transactional circuit into the asset circuit — is not imposed here as a behavioral stipulation. It is derived as the steady-state outcome of a worker optimizing lifetime consumption when the framework itself already provides their equity accumulation through the locked Stable Floor. The derivation also makes explicit where (A3) fails, which turns out to be the same immature-system window identified in Section 3.2.

5.1 Two-account households

A representative worker enters the period with labor income w_t (paid in transactional dollars, i.e. into *M^T*) and a locked Stable Floor balance F_t (equity, in the asset circuit *M^A*), which the worker cannot access before age T = 65. The floor accumulates exogenously to the worker's choices: K1 and K2 deposit into it on the schedule of Section 4, and it compounds at the equity return. The worker chooses consumption c_t and *discretionary* private saving *b_t* — voluntary purchases of equity beyond the floor — to maximize lifetime utility subject to the labor-income budget:

c_t + *b_t* = w_t + *y_t*,

where *y_t* is any spendable transfer (the KI dividend in Mode C; zero in Modes A and B). The flow that matters for Proposition 1 is *b_t*: it is precisely the worker-driven leak from *M^T* into *M^A* that (A3) requires to be zero on net.

5.2 Result 2 (the no-net-saving condition)

Claim. If the worker's desired lifetime equity exposure is *E** and the locked floor is on track to deliver retirement equity of *F_T* ≥ *E**, then optimal discretionary saving is *b_t* = 0, and the net worker flow from *M^T* into *M^A* is zero. Assumption (A3) holds in steady state.

Derivation. The worker's retirement resources are the sum of two equity stocks: the locked floor *F_T* and any discretionary holdings accumulated from the *b_t* stream. A life-cycle optimizer targets a total desired equity position E^* set by the consumption-smoothing problem (the standard permanent-income target). Discretionary saving is the residual needed to reach that target:

$b_t > 0$ only if the projected floor F_T falls short of E^*.

Under the framework's mature calibration, the floor is large by construction: the architectural paper's own figures put F_T at approximately $413,000 in Mode B (and $233,000 in the deflationary Mode A) on the realizable basis (Section 6.7), in launch-year purchasing power — at or above the equity position a median worker would otherwise target. When $F_T \geq E^*$, the optimizer has no reason to add discretionary equity: the marginal retirement-saving motive is already satisfied by the system. So $b_t = 0$, labor income is fully consumed in the transactional circuit ($c_t = w_t + y_t$), and no worker-driven flow crosses from M^T into M^A. This is Result 2, and it is exactly assumption (A3).

The mechanism is worth stating plainly because it is the non-obvious part: (A3) does not require workers to be forbidden from saving, nor does it assume they are hand-to-mouth. It holds because the framework *pre-supplies the saving motive that would otherwise drive the leak.* The locked floor crowds out the discretionary $M^T \to M^A$ flow not by prohibition but by satisfaction of the underlying objective.

5.3 When Result 2 fails — and why it is the same window as Section 3.2

The condition $F_T \geq E^*$ is not automatic. It fails in two cases, both bounded and both informative:

1. **The immature system (early transition).** Workers alive at launch, or born shortly after, have floors that have not yet compounded to E^*; the oldest transition cohorts may never reach it. For these workers $b_t > 0$ — they top up discretionary equity, producing a genuine $M^T \to M^A$ flow. This is the *same* immature-system window flagged in Section 3.2, where the locked-float share s_t is small: when floors are young, both the price-leak channel (3.2) and the worker-saving channel (5.2) are at their largest, and they shrink together as the system matures ($s_t \to 1$, $F_T \to E^*$ for successive cohorts). The framework's price-stability and distributional clean-ness are both weakest at launch and strengthen over the same horizon — a single maturation story, not two.
2. **High-target workers.** A worker whose desired equity exposure exceeds even the mature floor ($E^* > F_T$) saves the difference discretionarily regardless of system maturity. This produces a positive but *bounded* $M^T \to M^A$ flow concentrated among high-savers. It does not threaten Proposition 1 in aggregate because (a) it is a level shift in M^A, not an ongoing source of M^T contraction once the target stock is reached, and (b) it is precisely offset in the price identity by being saving the worker would have done under any monetary regime — it is not induced by the issuance rule.

5.4 Distributional reading

Two distributional implications fall out of the same block. First, because the floor satisfies the median worker's equity-saving motive, the framework converts what is currently a regressive channel — equity ownership concentrated among those with disposable income to invest — into a universal one, without requiring workers to forgo consumption to participate. Second, the residual discretionary saving b_t that survives Result 2 is concentrated among high-target (typically higher-income) workers, so the *marginal* private asset circuit remains theirs; the framework universalizes the *base* equity stake without claiming to equalize all asset holdings. This is the honest scope of the distributional claim: a universal floor, not a flat distribution.

5.5 Labor supply and the wealth-effect objection

The objection. The framework hands every citizen a locked floor on the order of $413,000 — roughly six to seven times median annual labor income. The standard worry follows at once: wealth of that size produces an income effect on labor supply, so citizens work less, output falls, and because K2 issuance, the consumer dividend, and the demographic absorption result are all indexed to the growth rate g, a labor-induced fall in g could be self-undermining. The mechanism is real; what the objection misjudges is the magnitude, and the reason is the lock.

A flow, not a stock. A worker with Frisch elasticity ν who values consumption and dislikes hours sets the marginal value of the wage equal to the marginal disutility of work. With spendable non-labor income Y_F from the floor system, the labor ratio ℓ ≡ h/h⋆ (hours relative to the no-floor benchmark) solves ℓ^(1+1/ν) + b·ℓ^(1/ν) = 1, whose single sufficient statistic is b ≡ Y_F/(ω h⋆), the floor's spendable income as a fraction of baseline labor income. The response is bounded: dℓ/db = −ν/(1+ν) at the benchmark, so each percentage point of b costs at most ν/(1+ν) percent of labor — about one-third at ν = 0.5. What moves labor supply is the spendable flow b, never the floor stock directly.

Why the lock is load-bearing. If the floor were liquid — principal freely consumable — the relevant flow would be its full annuitized return, b_liq ≈ rF/(ω h⋆). For a floor of six to seven times labor income at r = 4.5% that is b ≈ 0.3 and a labor reduction near nine percent: in that world the objection has real force. The Citizens Standard does not live in that world. The principal is locked, so only a distributed flow Y_F = ρ_eff·F reaches the budget and b = (ρ_eff/r)·b_liq. Under the wealth-maximizing default κ_d = 0 the distributed share is nil and the residual is the bequest-net annuitized value alone; for plausible distribution settings ρ_eff sits near one percent and the labor reduction falls to one-to-two percent. The lock is not incidental packaging — it is the device that turns a wealth stock which would gut labor supply into a spendable flow that barely perturbs it.

A level, not a trend. Because the wage and the floor both scale with the economy along a balanced path, b is invariant to the level of that path: the floor shifts the level of labor and output once, it does not bend the trend growth rate. Trend g is unchanged in the neoclassical benchmark; in an endogenous-growth reading where aggregate effort feeds g with elasticity ε_g, the trend falls by at most ε_g·[ν/(1+ν)]·b — bounded, and not a spiral, precisely because b does not grow with the floor stock.

An automatic stabilizer. Growth-indexing of the distributed flow closes the loop in the stabilizing direction. A fall in g lowers the growth-indexed flow, which lowers b, which raises labor, which restores g: the labor–growth map has loop gain ε_g·[ν/(1+ν)]·(d ln b/d ln g), far below one for any relevant calibration (≈ 0.016 in the illustration of Figure 2), so the fixed point is locally stable and self-correcting. The marginal-product response reinforces it — thinner labor raises the wage, lowering b further. The same growth-indexing that disciplines prices thus also disciplines labor supply. Proposition 4 (Appendix A.2, proof A.7) states the four claims; ν and b are the empirical quantities (A.5).

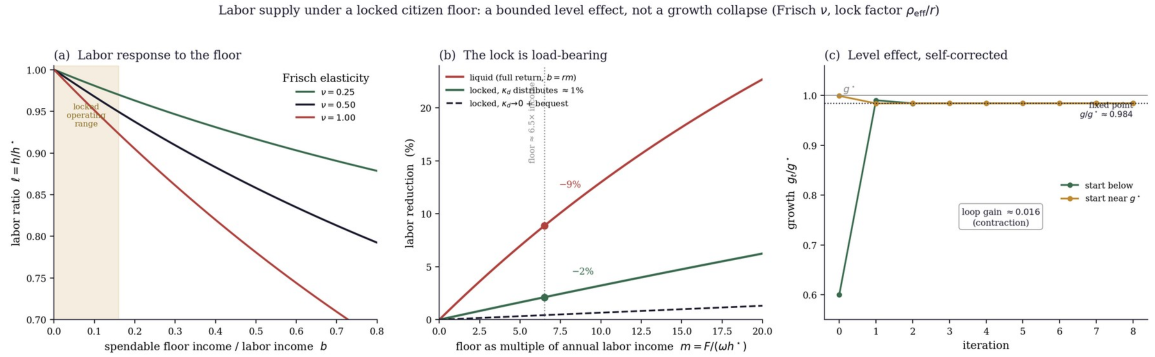

Figure 2. Labor supply under a locked citizen floor. (a) the labor ratio ℓ = h/h⋆ against the spendable floor-income ratio b, for Frisch elasticities ν; (b) the labor reduction against the floor's size as a multiple of labor income — a liquid

floor of ≈6–7× income would cut labor ≈9%, while locking the principal (distributing only ρ_eff·F) brings it to ≈1–2%; (c) with a growth-indexed flow the labor–growth loop is a contraction (gain ≈0.016), so growth settles at a bounded level ≈1.6% below g⋆ from any start. Parameters illustrative; ν and b are the empirical quantities of Appendix A.5.

5.6 The welfare-optimal dividend share

Wealth-maximizing is not welfare-maximizing. The framework treats κ_d = 0 — distribute no current dividend, let every growth-matched issuance compound in the locked floor — as the wealth-maximizing default, and it is: retaining the whole flow trivially maximizes the accumulated floor. But maximizing terminal wealth is not maximizing welfare. A citizen values consumption now as well as wealth later, and κ_d is exactly the dial between them: a share κ_d of each period's growth issuance G is paid as a current dividend (consumed in the working phase) and the remainder (1−κ_d)G is locked, compounded at r, and consumed in retirement or bequeathed.

That this maximizes the floor is an artifact of holding the return fixed. At universal scale the return is endogenous (§6.7): retaining the whole flow at full budget deepens the stock and pulls the floor's own return down, and beyond the feasibility bound the marginal retained dollar buys asset-price inflation rather than ownership. So even on pure realizable-wealth grounds the optimal κ_d is positive for the full-budget Mode — the feasibility condition and the welfare calculus below select the same direction, and the κ_d = 0 corner is available only to the low-budget Modes whose floor capture stays inside the bound.

A retention premium governs the split. With log utility, and with the lock keeping the labor wealth-effect tied to the current dividend alone (b = κ_d G, Proposition 4), the welfare-optimal share depends on a single sufficient statistic — the retention premium R = βΔ(1+r): the time-preference factor β, the terminal-value weight Δ (one if the floor is self-consumed in retirement, less than one if it is bequeathed and the bequest is weakly valued), and the gross return. The optimum κ_d* falls as R rises (Figure 6a): a higher return or more patience makes retention more attractive and distribution less, while impatience, a low return, or a mostly-bequeathed-and-weakly-valued floor all push κ_d* up.

When the default is optimal. The wealth-maximizing corner κ_d* = 0 is reached only when R exceeds about 1.2 — only under an effectively negative net discount, a planner weighting accumulation and heirs above current consumers. Under any conventional positive discounting the welfare-optimal dividend is strictly positive (Figure 6c): the κ_d = 0 default is revealed not as a neutral technical choice but as a deliberate accumulation-prioritizing value judgment, which the parameter makes explicit and tunable. The Proposition 4 labor distortion does not overturn this — by the envelope theorem it is second-order at the citizen's own labor optimum, entering κ_d* only through the budget, and becomes a first-order reason to lower κ_d only under a labor wedge such as a tax or externality. Proposition 8 (Appendix A.2; proof and Figure 6 in A.11) states the result; β, Δ, and r are the empirical and normative quantities (A.5).

6. Asset-market dynamics under a permanent structural buyer (Proposition 2)

The framework makes the FDCA a permanent, price-insensitive purchaser of total-market equity: every year, K1 and K2 dollars buy index shares regardless of valuation. The gross purchasing is permanent; the FDCA's net position, as Sections 6.1 and 6.6 show, turns from buyer to bounded seller as the floors mature. The natural objection — the second hard referee question, after the price-level question of Section 3 — is that a permanent inelastic buyer must bid prices up and *compress forward returns*, and that if equity returns compress toward the risk-free rate, the entire Stable Floor proposition (which assumes ≈ 7% nominal compounding) erodes over time. This section takes that objection seriously and locates where the return path settles. The result is that compression is a **bounded, one-time valuation level effect**, not a runaway return collapse, for two structural reasons the framework supplies by construction. It formalizes the intuition stated qualitatively in the architectural paper (Neo-Solon, 2026a, Section 12.7) and the intergenerational compression noted in the transition paper (Neo-Solon, 2026c).

6.1 The three forces

Let *f_t* be the tradable float (shares available to the market), *Q_t* an asset-valuation index (price relative to the replacement cost / earnings base of the underlying capital), and *r_t* the forward equity return. The float evolves as *Δf_t = I_t − A_t*: firm-side net issuance *I_t* adds to it, and net FDCA absorption *A_t* (defined below) subtracts from it. *V*aluation *Q_t* rises when the float shrinks relative to demand. Three forces act on this system:

1. **Float-shrink (compression).** Gross FDCA purchases move shares out of the float into locked accounts. With demand fixed, a smaller float raises *Q_t* and lowers forward *r_t*. This is the genuine compression channel and the heart of the objection.
2. **Drawdown (decompression).** The floors are not a sink. From age 65 each cohort draws its floor down at a sustainable rate, and at death the residual passes to an heir's floor rather than being sold; the resulting drawdown flow *L_t* (Section 3.5) *sells* shares back into the float. The structural buyer is also a structural seller. The relevant quantity for valuation is therefore **net** absorption, *A_t* = (gross purchases) − *L_t*, not gross purchases.
3. **The productivity anchor (Q-channel).** Long-run returns are not pure scarcity. They are tethered to the real return on productive capital — the earnings yield of the underlying firms. When *Q_t* rises (earnings yield falls), issuing new equity becomes cheap for firms relative to the return on real investment, so firm-side equity supply expands. New issuance enlarges the float, capping *Q_t*. This is the Tobin-Q mechanism: the FDCA can bid the *price* up, but it cannot permanently push the *return* below the productivity of capital without triggering a supply response that re-anchors it.

The structural-buyer objection has two distinct parts, and they are best kept apart. One is a price question: does a permanent, consistent buyer distort valuations? The other is a return question: can each citizen's floor realize the return its accumulation assumes when every citizen accumulates at once? Sections 6.2–6.6 settle the first — net absorption is bounded, the supply side responds, and the valuation premium is small. Section 6.7 settles the second. Under primary issuance, where floor funding builds new productive capital rather than competing for a fixed stock of existing shares, the return the floor earns is the marginal product of that capital, and it is endogenous: it falls as the citizen stake deepens the stock. The two analyses are complementary — one bounds the price distortion around the return, the other derives the return itself — and §6.2's anchoring of r* to the real return on capital is exactly the quantity §6.7 then solves for.

6.2 Proposition 2 (bounded compression)

Claim. Under the framework's issuance rule, forward equity returns settle at a bounded valuation premium above the historical baseline — a one-time level rise in `Q` — rather than a secular decline toward the risk-free rate. Specifically:

Net absorption *A_t* is positive while the floors accumulate and turns negative at maturity, as drawdown sales overtake purchases (Section 6.6). The steady state is therefore not a near-balance but a bounded structural-seller flow, so *A** stays bounded in both phases rather than tracking gross issuance. Compression is governed by *A**, not by gross purchases: it builds modestly during accumulation and is relieved — not deepened — as the floors mature and the buyer becomes a bounded net seller, with the bound set by the r − g rebalancing of Section 6.6 and by the Q-channel below.

- The **Q-channel** caps the valuation premium: Q_t rises only until the earnings yield falls to the point where net new equity issuance equals residual net absorption. At that fixed point $Q^*$is finite and r^* is anchored to the real return on capital, not to scarcity.

Derivation sketch. Write valuation dynamics as $\Delta Q_t = \theta \cdot (A_t - I_t)$, where I_t is net new firm equity issuance and $\theta > 0$ maps excess net demand into valuation change. Issuance responds to valuation through the Q-channel, $I_t = \varphi \cdot (Q_t - Q_{baseline})$, with $\varphi > 0$: when assets are priced above the replacement/earnings baseline, firms issue. The fixed point $\Delta Q = 0$ requires $A^* = I^* = \varphi \cdot (Q^* - Q_{baseline})$, so

$$Q^* = Q_{baseline} + A^*/\varphi.$$

The valuation premium is A^*/φ — finite, increasing in net absorption A^*, decreasing in the elasticity of equity supply φ. Forward return at the fixed point is the earnings yield at Q^*, which sits *below* baseline by the premium but *above* the risk-free rate so long as A^* is finite and $\varphi > 0$. There is no path to return collapse unless equity supply is perfectly inelastic ($\varphi \to 0$) **and** net absorption is permanently large (A^* does not fall with demographic maturity) — and the framework's own drawdown structure works against the second condition.

6.3 The honest concession: the accumulation phase

The clean steady-state result hides a real transitional effect, and it should be stated rather than buried. Before the age distribution matures — the same accumulation window as Sections 3.2 and 5.3 — there are many depositing young cohorts and few drawing-down retirees, so net absorption A_t is positive and larger than its steady-state value. During this phase a genuine valuation premium builds: Q_t rises and forward returns run modestly below the historical baseline. This is the formal content of the "modest intergenerational return-compression effect" the transition paper notes.

Two things make this concession survivable rather than fatal. First, it is *bounded and self-limiting* — the Q-channel caps it, and it eases as drawdown sales ramp up with cohort maturity. Second, and characteristically for this framework, the premium accrues to the floors as capital gains: the citizens holding the appreciating equity are the same citizens the system is accumulating wealth for. Early cohorts experience lower forward yields but higher realized valuations on the way there; the effect transfers value across the valuation path rather than destroying it. As in Sections 3.2 and 5.3, the framework's most strained window is its youth, and the strain resolves on the same maturation horizon.

The incidence wrinkle at launch (who captures the early premium). There is a distributional subtlety inside the accumulation-phase concession that should be stated plainly rather than smoothed over, because it is the sharpest version of the objection. At steady state the appreciation accrues to citizen floors — but at launch those floors are nearly empty, so the *early* premium accrues disproportionately to whoever already holds equity, which is the existing, already-wealthy asset distribution. In the first years, the structural buyer bids up a market owned mostly by incumbents, and citizen floors capture a growing share of the appreciation only as they accumulate. The benefit broadens from "existing holders" to "all citizens" over the maturation horizon rather than at the outset.

Three things bound this honestly. First, the flow is gradual: net FDCA purchases run on the order of 0.39% of market capitalization per year (§6.5), so the incumbent windfall in any single launch year is a fraction of one percent of valuation, not a step-change — there is no large one-time transfer to existing holders, only a slow appreciation they share with a growing citizen float. Second, the crossover is datable. Citizen floors hold share s_t of total equity; the appreciation captured by citizens versus incumbents is, to first order, in proportion $s_t : (1 - s_t)$. Citizens capture the *majority* of each year's premium once $s_t > 0.5$. On the locked-float trajectory used throughout this paper, s_t reaches one-half roughly two to three decades into the program (the same horizon on which the price-leak of §3.2 and the saving channel of §5.2 resolve), and approaches unity thereafter — so the regime in which incumbents capture most of the premium is the first ~20–30 years, declining monotonically, not a permanent feature. Third, the incumbent capture is of a *premium the framework itself creates*; absent the program there is no premium to capture, so the counterfactual is not “citizens get it instead” but “the premium does not exist.” The honest claim is therefore narrow: the framework front-loads a modest, gradual valuation premium that early-on accrues mostly to existing holders and progressively to citizens, crossing into majority-citizen capture within the accumulation window. It is a transitional incidence effect on a self-resolving horizon, not a structural transfer to the wealthy — but it is real, and a launch-era critic who holds equity is, for a time, a beneficiary.

The consumer-dividend channel (Section 4.3) acts directly on this concession. Because K3 pays a share κ_d of the growth budget to citizens instead of using it to purchase index shares, it lowers gross FDCA absorption — and therefore the very premium incumbents capture — by construction: a positive κ_d both shortens the window in which incumbents are net beneficiaries and shrinks its magnitude. The lever that delivers citizens current income is the same one that drains the launch-era premium, so the incidence wrinkle above is most acute precisely in the κ_d = 0 configuration and is dialed down monotonically as the dividend is turned up.

6.4 Consistency with the series baseline

This is why the architectural paper’s use of a 7% nominal baseline (Neo-Solon, 2026a, Section 12.7) is defensible rather than optimistic. Proposition 2 says the framework produces a bounded valuation premium with forward returns somewhat below a pure-scarcity-free baseline during accumulation, re-anchoring to capital productivity in steady state. That is consistent with the architectural paper’s two-sided position — that curtailed monetary expansion argues for somewhat lower returns while direct, consistent equity demand argues for neutrality — and with the empirical paper’s strategy of stress-testing the whole return distribution rather than betting on the mean (Neo-Solon, 2026b). Proposition 2 does not claim to pin down r^* to a number; it claims the weaker, defensible thing — that the structural-buyer objection does not imply return collapse, because the buyer is also a seller and the supply side responds.

6.5 Illustrative magnitude of the premium

As with the price leak, an order of magnitude can be put on the valuation premium using figures already on record, with the same caveat: this is a *worked illustration, not a calibration*. The gross-purchase figure is anchored; the drawdown fraction, total equity capitalization, and supply elasticity are bracketed within plausible ranges and are the quantities a full empirical exercise would estimate. The point is to show the premium is small under reasonable assumptions and to expose the inputs so a reader can recompute.

Table 2. Structural-buyer premium calibration inputs (Section 6.5).

Input	Symbol	Value used	Status
Gross annual FDCA equity purchase	—	≈ \$272B/yr	Anchored (Neo-Solon 2026a)
US total equity capitalization	—	≈ \$69,000B	Wilshire 5000 full-cap, YE 2025 (~\$69T, order of magnitude)
Drawdown at demographic maturity	L_t	40%–90% of gross	Bracketed
Normalized equity-	φ	0.5–2.0	Assumed

Input	Symbol	Value used	Status
supply elasticity			

Net absorption at demographic maturity is $A^* = \text{gross} - L_t$. As retiring cohorts' drawdown ramps up, L_t rises and A^* falls:

- At L_t = 40% of gross: A^* ≈ \$163B/yr ≈ 0.24% of equity capitalization.
- At L_t = 60%: A^* ≈ \$109B/yr ≈ 0.16%.
- At L_t = 90%: A^* ≈ \$27B/yr ≈ 0.04%.

The steady-state valuation premium is A^*/φ expressed against the float. Taking the mid drawdown (A^* ≈ \$109B, ≈ 0.16% of capitalization) and varying the supply elasticity:

- $\varphi = 0.5$ (inelastic supply): premium ≈ 0.32% on the valuation level.
- $\varphi = 1.0$: premium ≈ 0.16%.
- $\varphi = 2.0$ (elastic supply): premium ≈ 0.08%.

A one-time valuation premium of this size raises price-to-earnings by the same percentage and lowers the forward earnings yield by roughly that fraction of its level — for example, a 0.5% premium on a 5% earnings yield trims the forward yield by about 0.025 percentage points. The effect is a one-time level shift, not a cumulative drag: it does not compound year over year, and it eases as A^* falls with demographic maturity. The reading: across the plausible range, the structural-buyer valuation premium is on the order of a few tenths of one percent of valuation and a small fraction of a percentage point of forward yield — material enough to acknowledge, far too small to constitute the return collapse the objection envisions. A full calibration would pin L_t, total capitalization, and φ to data; the conclusion that the premium is small is robust across the bracket.

6.6 The maturity question: a bounded rebalancing, not a sell-off

A natural objection to a universal floor is temporal: once the first cohorts reach 65 and begin drawing down, does the accumulated stock reverse into a wave of selling that floods the market? The honest answer bounds the selling rather than denying it. Two features rule out the crash the objection pictures. First, the flow is gradual: cohorts reach 65 one year at a time and draw income at a sustainable rate (the 4% convention), so the onset is a slow ramp, not a cliff. Second, a decedent's residual balance does not leave the equity circuit at all — under the inheritance rule it passes into a beneficiary's own Stable Floor under the same lock and continues to compound, so there is no liquidation spike at death.

What remains is a genuine, gradual shift in the floors' net flow, and it is the dated form of the structural-buyer/seller result already stated in Section 6.2. The price-pressure-relevant quantity is net external absorption — new money issued into floors (buying) minus the consumption retirees withdraw (selling); reinvested dividends are not independent demand, since in aggregate a reinvested dividend buys back exactly the ex-dividend price it is funded from. An overlapping-generations model of this flow (Appendix A.6) shows the floors are net buyers through the accumulation decades and cross into net sellers around the time the first cohorts have been retired for a decade — near year 55 at the central calibration. The buyer phase is the first half-century; the structural-seller phase follows. This is the same transition Proposition 2 names; the model simply dates it.

The structural-seller phase is bounded, not a flood. Because the locked floor earns a return r above the economy's growth rate g, the floors' holdings would otherwise outgrow the economy without limit — the $r > g$ dynamic, here operating on a universally held base. The steady-state net outflow is the rebalancing of exactly that excess: at a bounded market share it equals $(r - g)$ times the floor stock, on the order of a few percent of GDP per year, smooth and predictable. It is the same flow the Q-channel absorbs through firm-side issuance in Proposition 2, and κ_d

moderates it directly by routing part of the growth budget to consumption rather than into floors (Section 6.3). When r does not exceed g the excess and the structural selling vanish entirely. The inheritance rule earns its place on price-stability grounds, not only fairness. The same model with floors liquidated at death rather than inherited adds a death-liquidation spike on top of the ongoing withdrawal flow — roughly doubling the mature outflow and making it lumpy. That floors are inherited under the same lock converts a potential liquidation shock into the smooth, bounded rebalancing above. The honest claim is therefore narrow: the framework rules out a maturity crash and bounds the steady-state outflow to the r − g rebalancing, but it does not make the floors permanent net buyers. At maturity they are net sellers of a bounded, absorbable magnitude, and whether that magnitude is absorbed smoothly is the same concentration question Proposition 2 already carries — governed by the Q-channel and κ_d — not a sudden flood.

6.7 The realizable floor: a general-equilibrium account

Sections 6.2–6.6 hold the return on capital fixed and ask how far the structural buyer moves prices around it. This section solves for the return itself, because the floor figures depend on it and because a universal program cannot take its own return as given.

Why the return is endogenous. A citizen's floor is built by compounding deposits at the return on capital. For one citizen considered in isolation, that return is exogenous — the price-taker return a marginal participant earns by buying into a market priced by everyone else. For the program as a whole it cannot be, because the program is not a marginal participant but a universal one, and a figure that holds for one citizen need not hold for all citizens at once. The point is concrete: compounding the full-rate deposit at a fixed market return across all citizens, with no drawdown, the aggregate floor would reach on the order of 460% of the entire U.S. equity market. That is not an ownership share — nothing can be owned several times over — but the signature of a fixed return applied at a scale where the return is no longer fixed. At universal scale the return has to be solved in general equilibrium.

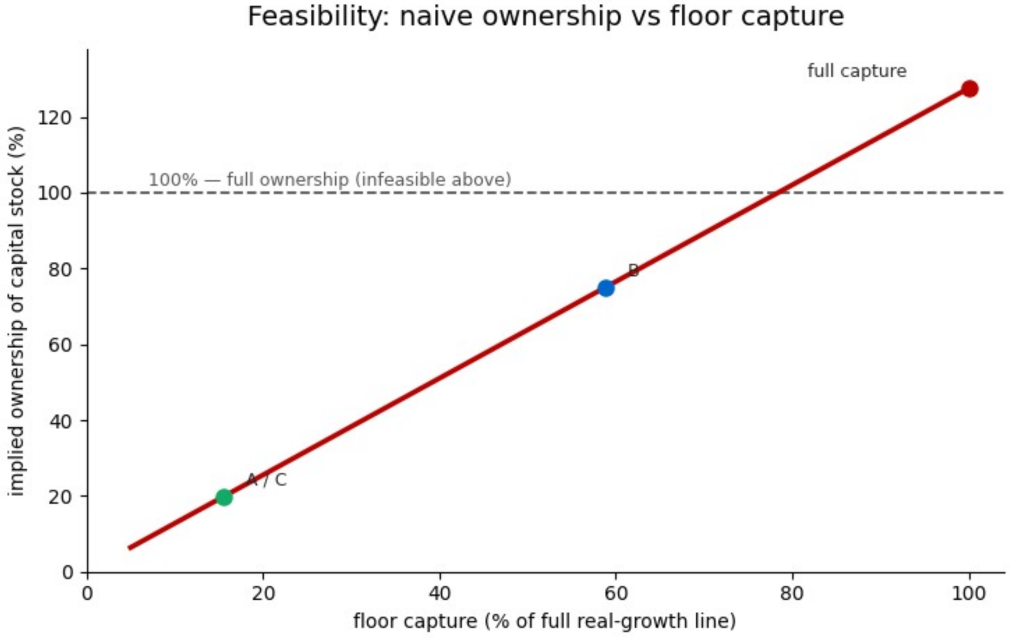

Figure 8. Implied ownership of the capital stock rises with floor capture; full capture (the unrecalibrated large-floor configuration) crosses 100%, while the recalibrated 60/40 split sits safely inside the bound.

The feasibility condition. This yields a calibration constraint that holds with the force of a conservation law. The K1 and K2 floor channels purchase ownership of real productive capital, so a calibration is feasible only if the aggregate accumulated stake stays within the capital that exists — at most 100% of the investable stock, and in practice a target share well below. Issuance directed into the floor beyond what the stock can absorb as genuine new capital does not create wealth; it bids up the price of existing claims, and sustained excess compounds that into asset-price inflation — a bubble carried as “floor.” Issuance beyond the capital channel’s absorptive bound therefore belongs in the dividend channels (K3, KI), which distribute consumption and carry no ownership ceiling. Two inflation channels follow, and the issuance envelope divides between them: over-directing K1/K2 produces asset-price inflation; over-directing K3/KI produces consumer-price inflation. The real-growth-matched envelope fixes the total; the feasibility condition governs the split.

The bound binds in two forms. The hard form is the ownership impossibility above. The soft form, which binds first, is endogenous: as the citizen stake deepens the stock, the marginal product of capital falls, so the return on the marginal floor dollar declines, and once it reaches the social discount rate an added dollar is worth more as current income than as locked floor. A feasible calibration fills the floor to that point and dividends the remainder.

Solving for the floor. Under primary issuance the citizens come to own a real, growing share of the productive capital stock, and the relevant return is that capital’s marginal product. A standard Solow/overlapping-generations calibration (capital share $\alpha \approx 0.35$, depreciation $\delta \approx 0.05$, baseline $K/Y = 3$, real growth 2%) places the no-program net return on capital at 6.67%. As the program runs, the floor reinvests its return and accumulates, the citizen stock deepens the aggregate stock, and the marginal product falls. The deepening does not run away: the return falling toward the growth rate slows accumulation and finite lifecycles cap each cohort’s holding, so the citizen share reaches a stable steady state. The return the floor earns is this attenuated, with-program marginal product. The same $r > g$ force that in private hands concentrates capital is here harnessed and self-limited — broad ownership compounds at r until that very breadth pulls r toward g.

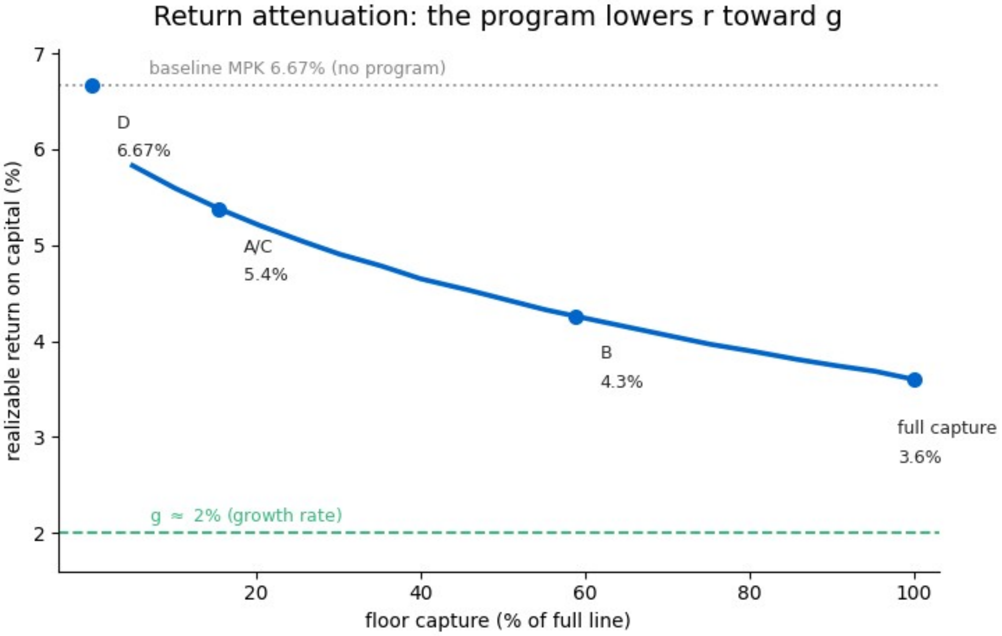

Figure 9. The realizable return on capital falls from the 6.67% no-program baseline toward the growth rate as floor capture rises; each Mode sits at its capture rate — Mode D, building no floor, sits at 0% capture and earns the full unattenuated 6.67% baseline.

The floor by Mode. The attenuation scales with how much of the real-growth line a Mode directs into the floor — its capture rate. Conservative Modes barely move the stock; the balanced Mode moves it materially. Reading the floor off the engine at each Mode's realizable return:

Table 3. Realizable Stable Floor and dividend by Mode.

Mode	Floor capture	Capital share	Realiz. return	Realizable floor	Dividend
A (deflationary)	15.5%	~16%	5.38%	≈ $233K	— (+ deflation gain)
B (60% floor / 40% K3)	58.8%	~25%	4.26%	≈ $413K	$513/yr K3 ($57K lifetime)
C (low floor + KI)	15.5%	~16%	5.38%	≈ $230K	$1,293/yr KI ($143K lifetime)
D (distributed)	0%	~0%	6.67%	$0	$672/yr dividend ($75K lifetime)

Real 2025 dollars. The floor is a heritable, locked capital stake, not a retirement balance: the citizen lives on its yield and bequeaths the principal. Figures are general-equilibrium estimates from a stylized model, reported as such. Mode D builds no locked floor: its benefit is delivered entirely as the standing dividend (the largest of any base Mode, the full g_r·M^T leash, ≈ $672/yr at launch), and the 6.67% figure is the unattenuated return on capital that a citizen earns by privately reinvesting the dividend rather than a return on a system-built floor — Mode D, deepening no aggregate capital stock, leaves the fundamental return uncompressed.

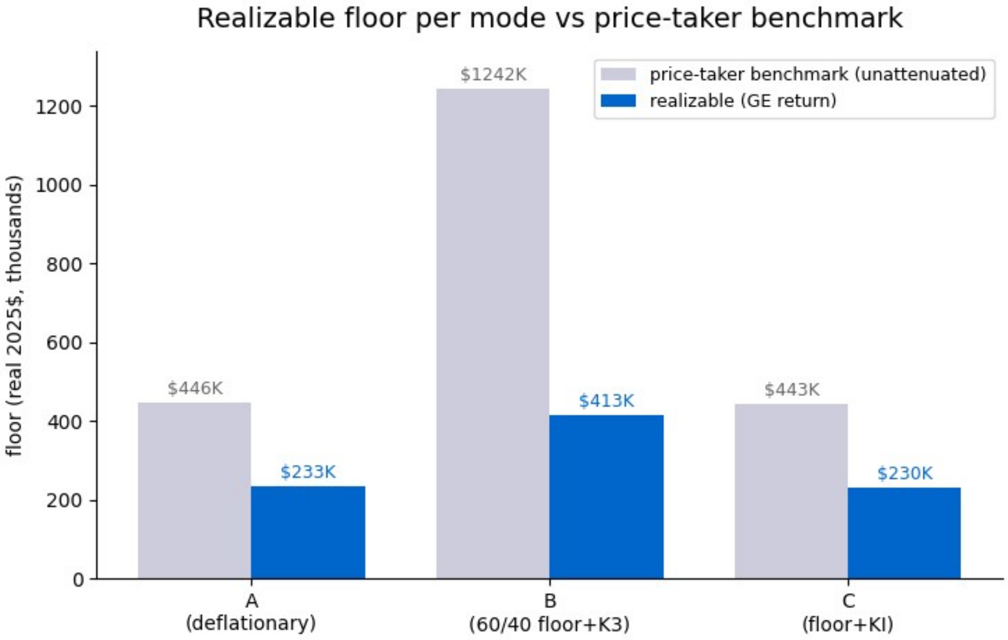

Figure 10. Realizable floor by Mode (at the attenuated return) against the price-taker benchmark (at the unattenuated marginal product). The gap is the attenuation, largest where capture is highest.

Two features stand out. The conservative Modes A and C earn ≈5.4% — above the price-taker return a single participant would assume — because they capture so little of the line that they barely attenuate the fundamental return; their floors are robust. And the balanced Mode B directs 60% of the line to the locked floor and 40% to the K3 consumer dividend. This keeps its defining property — Mode B remains approximately price-stable at the balance point, because K3 reallocates within the issuance budget rather than adding to it (derived CPI = 0.00%) — while sitting comfortably inside the feasibility bound: capture 58.8%, implied ownership 75%, a realizable floor near $413K, and a $513/yr K3 stream. Mode B is the largest-floor option and a genuinely pickable one. The taxonomy is four Modes — A deflationary, B a balanced large floor plus dividend, C a small floor plus current income, and D an all-dividend configuration with no floor — spanning the feasible range of the floor/dividend split, from a pure floor (A) to a pure dividend (D).

The compression is the result, not the cost. The attenuated return is not damage to be minimized. The program drives the economy-wide return on capital from 6.67% toward the growth rate — to 4.26% in Mode B at a ~25% citizen capital share — and that compression of r − g is a central equality result. The mechanism that in private hands makes capital income outrun growth and concentrate wealth is here run in reverse: universal ownership narrows the gap between the return on capital and the growth of the economy precisely by broadening who holds the capital. Incumbent capital earns less; new entrants and labor gain. The floor is smaller than a single-participant projection for exactly the reason the framework works — because it is universal, and universality bends the return.

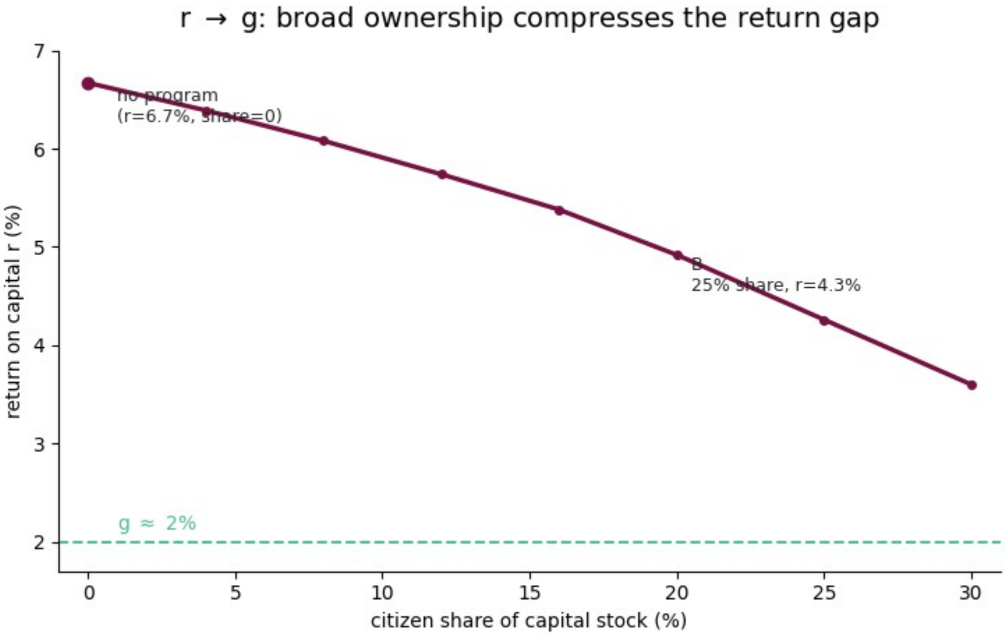

Figure 11. Broad ownership compresses r − g: as the citizen share of the capital stock rises, the return on capital falls from 6.67% toward the 2% growth rate.

Modeling note: the realizable figures derive from a stylized Cobb-Douglas Solow/OLG calibration on a domestic, closed-economy capital stock — the appropriate case for a state-owned national index funded by primary issuance. They establish order of magnitude and the three mechanisms — bounded ownership share, endogenous return attenuation, r − g compression — robustly across the α, δ range tested. A fuller calibration would tighten the absolute floors; the qualitative conclusion holds across every variant.

7. Crisis-mode operations and automatic stabilizers

The preceding sections establish the framework's steady-state and accumulation-phase behavior. This section addresses the question a monetary economist will ask first: how does it respond to shocks? The answer must be supplied carefully, because the framework has deliberately given up the conventional stabilization channel. By fully reserving the payment system and separating it from credit (Neo-Solon, 2026a, Section 9), the Citizens Standard weakens the interest-rate transmission mechanism that a conventional central bank relies on. It therefore cannot lean on the standard tool and must replace it. The replacement has two tiers: a formula-based automatic stabilizer (KI under price-level-path targeting) and a discretionary backstop (the bounded, rules-triggered emergency toolkit of Neo-Solon, 2026a, Section 10). Section 3.5's slow/fast lever distinction is what makes this work, and it also bounds what the framework can do.

7.1 The stabilization assignment

It is worth stating the division of labor explicitly, because conflating the channels is the source of the most common objection (that K2 is procyclical). Three instruments, three roles:

- **K1/K2 — not stabilization instruments.** These are slow levers (Section 3.5): they feed locked floors and reach the transactional circuit only through the liquidation flow L_t on a decade lag. K2 tracks measured growth, so it issues *less* in downturns — procyclical if misread as demand management. The framework's response, following the architectural paper (Neo-Solon, 2026a, Section 2.3), is not to fix K2 but to relieve it of stabilization duty entirely. K2 is a long-horizon ownership channel; a single slow year is immaterial to a floor compounding over a working life.
- **KI — the automatic stabilizer.** The fast lever, entering M^T immediately and equally. Under price-level-path targeting it responds to the price-level gap without any committee decision. This is the framework's formula-based countercyclical instrument.
- **The emergency toolkit — the discretionary backstop.** Bounded, rules-triggered, automatically reversing tools (Neo-Solon, 2026a, Section 10) for shocks beyond what the KI formula absorbs.

7.2 Demand collapse — the automatic stabilizer engages

Consider a negative demand shock: spending falls, output gaps open, and actual CPI drops below the target path $(1 + \pi^*)^t$. The KI formula from Section 4.4,

$$\mathrm{KI}(t) = \max(0,\ \lambda \cdot \mathrm{gap}(t) \cdot M^T + [\pi^* + g_r] \cdot M^T - K3_{agg} - L_t), \quad \mathrm{gap}(t) = \ln\left[\, (1+\pi^*)^{(t-1)} / \mathrm{actual\_price\_level}(t-1) \,\right],$$

responds automatically. As `actual_CPI(t−1)` falls below the desired path, the ratio inside the logarithm rises, `gap(t)` increases, and KI issuance expands — depositing spendable dollars directly into the transactional circuit, equally across all citizens, with no intermediating decision. This is an automatic stabilizer in the strict Friedman sense: the rule itself leans against the shock. Two properties distinguish it from conventional stimulus. First, it reaches households directly rather than through the credit channel, so it works even when the credit channel is impaired — precisely the 2008-style condition where rate cuts lose traction. Second, it is distributionally flat by construction, avoiding the Cantillon-ordered transmission of conventional injection.

In Modes A and B, where KI is dormant (`KI = 0`), this automatic channel is unavailable — which is the price-stability cost flagged in Section 3.5. A society running Mode B for its steady-state price stability (approximate away from the transaction-ratio balance point) holds weaker automatic countercyclical capacity and leans correspondingly more on the discretionary toolkit. This is a genuine Mode-level tradeoff, not a free lunch: the regime that is calmest in steady state is the one with the least automatic shock response.

7.3 Inflation surge — a real asymmetry

The symmetric case is weaker, and the asymmetry should be stated rather than hidden. Under an overheating shock, actual CPI runs above the target path, the gap-closure term shrinks, and KI contracts. But KI is bounded below by zero: `max(0, …)`. The formula can stop adding money; it cannot claw back dividends already distributed. So the framework's *automatic* response to inflation saturates at "halt KI," whereas its automatic response to deflation ("expand KI") has no comparable ceiling. Above-target inflation beyond what halting KI absorbs must be met by the discretionary toolkit — reserve drains and the other Section 10 inflation-surge tools — not by the formula.

This asymmetry is not unique to the Citizens Standard (a conventional central bank also finds tightening into a downturn politically harder than easing), but it is sharper here because the spendable channel only runs one way. The honest characterization: the framework is a strong automatic stabilizer against deflation and a weak one against inflation, with the discretionary toolkit carrying the inflation-control load. A society that weights inflation risk heavily should note that this places more weight on the emergency tools functioning as designed.

7.4 Supply shocks — the same dilemma everyone faces

A negative supply shock — output down, prices up — is the hardest case for any monetary framework, and the Citizens Standard does not escape it. The price-level signal says contract; the output gap says expand; no single instrument resolves the conflict. What the framework can claim is narrow and defensible: because it targets a price-level *path* rather than the year-over-year inflation *rate*, and closes gaps gradually at speed $\lambda = 0.5$, a transient supply shock that bumps the price level does not trigger the violent contraction that strict inflation-rate targeting can produce. The path target treats a one-time level shock as a gap to be closed over several years, not an emergency to be stamped out immediately. This is a smoothing advantage, not a solution to the supply-shock dilemma — the framework still cannot simultaneously stabilize prices and output against a supply shock, because nothing can.

7.5 Comparison to a conventional reaction function

A conventional central bank follows a reaction function — a Taylor rule — setting one instrument (the policy rate) in response to the inflation and output gaps, transmitted through credit. The Citizens Standard differs on three axes. First, **instrument**: it acts through a direct, distributionally flat transfer (KI) rather than the price of credit, so transmission does not depend on a functioning credit channel or on the zero-lower-bound being slack. Second, **discretion**: the first-line response is formulaic and automatic, with discretion confined to the bounded emergency toolkit, where a Taylor-rule bank exercises ongoing discretion over the rate. Third, **scope**: the framework deliberately retains *less* fine-tuning capacity, having separated payment from credit; it trades the conventional bank's continuous discretionary control for rule-bound automaticity plus a bounded backstop. The framework's wager — consistent with the rules-versus-discretion tradition the architectural paper situates itself in (Neo-Solon, 2026a, Section 17.1) — is that a transparent automatic stabilizer plus a bounded backstop is more robust to capture and policy error than continuous discretion, at the cost of some fine-tuning capacity. Whether that wager is correct is an empirical and institutional question the model frames but does not settle.

7.6 A linearized dynamic system and impulse responses

From mechanisms to dynamics. The preceding subsections describe the system's response to shocks in words; this one writes it down. Collect the economy's deviations from the balanced path into three gaps — the price-path gap x_t of Section 3.6, an output gap $\tilde{y}_t$, and an asset-valuation gap v_t — and linearize: $x_t = (1-\psi\lambda)x_{t-1} + \kappa\,\tilde{y}_{t-1} + \lambda_{leak}\,v_{t-1} + e^{\pi}_t$; $\tilde{y}_t = \varphi_y\,\tilde{y}_{t-1} - \gamma\,x_{t-1} + (1-\omega_F)e^{y}_t$; and $v_t = \varphi_v\,v_{t-1} + e^{v}_t$. The first row is the Proposition 3′ recursion with two couplings added — an output-gap push κ and a small asset-circuit leak λ_{leak}. The system is stable at the baseline calibration (eigenvalues 0.12, 0.68, 0.90) and collapses to the pure gap recursion when the couplings vanish. The replication package simulates this system as a calibrated (not estimated) dynamic model: it reproduces these eigenvalues, computes the impulse responses to price, demand, and asset shocks — an asset-price shock reaches consumer prices only through the bounded leak, with a peak price-gap impact of 0.030 against the theoretical bound of 0.033 — and runs a randomized dynamic-robustness test over ten thousand admissible calibrations in which stability and circuit separation hold in every draw. The dynamics are therefore not an artifact of the baseline.

Demand collapse: the floor cushions. A negative demand shock e^y enters the output gap scaled by (1−ω_F), where ω_F is the share of household income drawn from the acyclical floor flow. Because that share does not fall in a downturn, the output trough shrinks linearly in ω_F — fifteen percent shallower at ω_F = 0.15 (Figure 4a) — while the price-path gap barely moves, since KI holds the anchor and the two-circuit structure keeps the demand → consumer-price slope κ small. This is the Section 7.2 automatic stabilizer made quantitative: a stock-based, acyclical income floor stabilizes demand with no discretionary act.

Cost-push: KI self-corrects. A cost-push shock e^π lifts the price-path gap; KI returns it at the rate set by its eigenvalue, ≈ 1−ψλ, so the gap is gone within a handful of periods (Figure 4b). Switch KI off (ψλ = 0) and the same shock leaves a unit root — the gap persists indefinitely. The throttle is what supplies the nominal anchor: the Section 7.3 point that the system owns a genuine corrective instrument, not a hope that shocks fade.

Asset shocks: contained by construction. A shock to the asset circuit e^v moves valuations fully but reaches consumer prices only through the leak: the steady-state pass-through is bounded by λ_leak/(ψλ), about three percent at the baseline (Figure 4c). The circuit separation that does the work in Section 3 is therefore also a containment property — asset-market turbulence is quarantined from the consumer price path. Proposition 6 (Appendix A.2; proof and Figure 4 in A.9) states the four results; κ, γ, φ_y, φ_v, λ_leak, ω_F are the empirical quantities a calibration would pin down (A.5).

8. Open problems

1. **Calibration of the steady-state circuit-tracking condition.** Establishing empirically that M^T can be held to Y growth across the cycle.
2. **The production-side fixed point (Proposition 2 — derived; illustratively bounded; full calibration outstanding).** Section 6 formalizes return dynamics under a permanent structural buyer: compression is a bounded valuation premium $Q^* = Q\_baseline + A^*/\varphi$, capped by the Q-channel and limited by demographic-steady-state net absorption A^*, not a secular return collapse. Section 6.5 provides a worked illustrative bound — a premium on the order of a few tenths of one percent of valuation across the plausible parameter range. What remains is a full empirical calibration of net absorption A^* (gross issuance net of the drawdown flow L_t at demographic maturity), total equity capitalization, and the equity-supply elasticity φ, against data. The conceptual objection is answered and the order of magnitude is bounded; a dated, data-grounded forward-return path is the empirical paper's work. The production-side fixed point and the seller-rebalancing (net-absorption) channel are taken up and closed in the companion structural-buyer paper (Neo-Solon, 2026h).
3. **The seller-rebalancing channel (derived; illustratively bounded; full calibration outstanding).** Section 3.2 turns the old assumption into a portfolio-choice result: the direct MPC out of a wealth-neutral swap is zero, and the residual leak into goods prices is $\kappa\_W \cdot (\Delta/\mu^*) \cdot (1 - s\_t)$, shrinking as floors mature. Section 3.2a provides a worked illustrative bound — a leak below roughly 1% of M2 across the plausible parameter range and near 0.4% in the central case at launch (Mode B), falling toward zero as floors mature. What remains is a full empirical calibration of the target money share μ^* and a dated path of the locked-float share s_t. This is no longer an axiom a skeptic can simply reject; the order of magnitude is bounded, and the precise early-transition leak is the empirical paper's work.

4. **Index-definition governance under a permanent structural buyer (a governance problem, not a calibration one).** The results above concern the *macroeconomics* of a permanent total-market buyer; they do not address the *political economy* of the index it tracks. Because inclusion in the tracked index confers guaranteed, price-insensitive, non-selling demand, index membership acquires a value it does not have under ordinary float-weighted indexing, and the body that defines inclusion criteria holds a corresponding concentration of power. Rules-based tracking of a pre-existing independent total-market index (CRSP-style) removes discretionary security *selection* — there is no “which companies” decision for the issuing authority to corrupt — but it relocates the leverage rather than eliminating it: whoever sets the inclusion methodology, and whoever governs that body, now sits at the capture point, and the stakes are higher than for any existing index fund precisely because the buyer never sells. This is distinct in kind from the three calibration problems above: it is not answered by estimating a parameter but by mechanism design — mechanical, exogenous inclusion criteria with no discretionary margin, and independent governance of the index methodology insulated from the issuing authority. Quantifying the inclusion premium (the valuation uplift conferred by guaranteed non-selling demand, which is the size of the capture incentive) and specifying a non-discretionary index-governance mechanism are open problems the present model does not close. They are flagged here so the governance surface is named rather than discovered. Both are addressed in the companion structural-buyer paper (Neo-Solon, 2026h), which specifies a mechanical, constitutionally-locked index-governance mechanism and quantifies the inclusion premium.

Technical Appendix

A.1 Notation

Symbol	Meaning	Introduced
`M_t`	Total money stock	3.1
M^T_t	Transactional circuit (prices goods)	3.1
M^A_t	Asset circuit (holds equity claims)	3.1
M^T_bank	Legacy bank-created money (transitory third pool)	3.4
`P_t`	Goods price level	3.1
V	Transactional velocity (stable)	3.1
Y_t	Real output	3.1
Δ	Size of an FDCA equity purchase (forced sale)	3.2
`m, e, W`	Seller’s money, equity, total wealth (`W = m + e`)	3.2
μ*	Target money share, μ* = m/W = μ(r, η)	3.2
κ_c	Consumption rate out of wealth, c = κ_c·W	3.2
κ_W	MPC out of asset wealth (≈ 0.02–0.05)	3.2
s_t	Share of risky assets held in locked floors	3.2
r, η	Equity return; convenience yield of money	3.2
L_t	Liquidation flow (matured floors → M^T)	3.3 / 3.5
α	K1 fraction of GDP per capita (= 0.025)	4.1
β_A	Mode A K2 fraction of real growth (= 0.125)	4.2
g_r, g_p	Real growth rate; population growth rate	4.1
π*	Inflation target (Mode C, = 0.02)	4.3
λ	KI gap-closure speed (= 0.5)	4.3
w_t, y_t, b_t	Labor income; spendable transfer; discretionary saving	5.1
F_t, F_T	Stable Floor balance; floor at	5.1

Symbol	Meaning	Introduced
	retirement	
E*	Worker’s desired lifetime equity exposure	5.2
f_t	Tradable equity float, *Δf_t* = *I_t* − *A_t*	6.1
Q_t, *Q_baseline*, Q*	Valuation index; baseline; fixed-point value	6.1 / 6.2
A_t, A*	Net FDCA absorption; steady-state value	6.1 / 6.2
I_t	Net firm equity issuance, *I_t* = *φ*·(*Q_t* − *Q_baseline*)	6.2
θ, *φ*	Valuation-adjustment speed; equity-supply elasticity	6.2
κ_d	K3 consumer share of the growth budget (Mode parameter)	4.4
K3_agg	Aggregate consumer dividend, = κ_d·(g_r·M2 − K1_agg); per citizen K3_agg / N	4.4

A.2 Propositions (formal statements)

Proposition 1 (two-pool price stability). Let (A1) M_t = M^T_t + M^A_t with P_t = M^T_t·V/Y_t and V constant; (A2) the seller-rebalancing result of A.3 below — issuance enters M^A with zero direct consumption response, so to first order it does not enter M^T; (A3) the no-net-saving condition of A.4 — workers make zero net discretionary transfer from M^T to M^A. Then if M^T_t/M^T_{t−1} = Y_t/Y_{t−1}, the goods price level P_t is constant, and issuance consistent with the rule affects asset values and floor balances but not the cost of living, up to the bounded leak of A.3. Spendable channels that enter M^T directly (KI, and the consumer dividend K3 of Section 4.3) are admissible under this Proposition only to the extent that L_t + KI + K3 does not exceed the Y-matched expansion of M^T; a K3 share κ_d that respects this ceiling preserves P_t, while one that breaches it is, by definition, a move off the price-stability locus into a Mode-C-type regime.

Proposition 2 (bounded compression). Let float evolve as *Δf_t* = *I_t* − *A_t*, valuation as *ΔQ*_t = *θ*·(*A_t* − *I_t*) with *θ* > 0, and firm issuance respond as *I_t* = *φ*·(*Q_t* − *Q_baseline*) with *φ* > 0. Then a fixed point exists at Q* = *Q_baseline* + A*/*φ* with A* = I*, finite for any finite steady-state net absorption A* and any *φ* > 0. Forward return at Q* lies below the no-absorption baseline by the premium A*/*φ* but strictly above the risk-free rate; return collapse requires both *φ* → 0 (inelastic supply) and A* non-decreasing in demographic maturity, the latter contradicted by the drawdown structure (*A_t* falls as retiring cohorts’ liquidation *L_t* rises).

Proposition 3 (convergence of the path-targeting dividend rule). Let *x_t* ≡ ln[target_price_level(t)/CPI(t)] be the log price-path gap, and let the KI rule of Section 4.4 inject g_KI(t) = *λ*·x_{t−1} (closure gain *λ*, baseline 0.5). With *ψ* > 0 the within-period pass-through from path-closure issuance to realized inflation, the gap evolves as *x_t* = (1 − *ψλ*)·x_{t−1} + ε_t. Then the rule converges without oscillation iff 0 < *ψλ* < 1; it converges with damped oscillation for 1 < *ψλ* < 2, and diverges for *ψλ* ≥ 2. Proof: the homogeneous solution is (1 − *ψλ*)^t x_0, stable iff |1

$-\psi\lambda| < 1$ (i.e. $0 < \psi\lambda < 2$) and non-oscillating iff $0 \le 1 - \psi\lambda < 1$ (i.e. $0 < \psi\lambda \le 1$). ∎ Since λ is a chosen gain, any measured ψ can be stabilized by selecting $\lambda < 1/\psi$; this is the sense in which path-targeting “self-corrects rather than oscillates” (Sections 3.5, 4.4) — a condition, not an assertion. The single-period pass-through is the conservative case: distributed-lag pass-through spreads the per-period response and enlarges the stable region. The Mode T damper KI_T (Section 4.7) has effective gain $\chi\lambda \le \lambda$, so it inherits convergence with strictly more margin and cannot oscillate on any setting the steady-state rule does not.

Proposition 3′ (structural pass-through and the maturing-circuit stability condition). Under (A1) $P_t = M^T_t \cdot V/Y_t$ with V constant, and a KI gap-closure rule injecting $\lambda \cdot x_{t-1} \cdot B$ dollars into M^T against base B, the path-gap obeys $x_t = (1 - \psi\lambda)x_{t-1} + \varepsilon_t$ with structural pass-through $\psi = B/M^T$. The feedback is locally stable iff $0 < \psi\lambda < 2$ and monotone iff $\psi\lambda < 1$ — equivalently $\lambda < 2M^T/B$ and $\lambda < M^T/B$. For broad-money indexing ($B = M2$), $\psi = M2/M^T = 1 + M^A/M^T$, increasing in system maturity, so any fixed λ determinate at launch is eventually violated; transactional indexing ($B = M^T$) gives $\psi \equiv 1$ and the maturity-invariant condition $\lambda < 1$, while an adaptive gain $\lambda_t = c \cdot M^T_t/M2_t$ ($c < 1$) holds $\psi\lambda = c$ at all dates. The predetermined liquidation flow L_t enters as forcing and contributes no root; stability depends only on the contemporaneous feedback gain. Under i.i.d. innovations of variance σ^2, the stationary gap variance is $\sigma^2/[1 - (1 - \psi\lambda)^2]$, minimized at $\psi\lambda = 1$; the gain is therefore set just inside this boundary, trading negligible variance for one-sided robustness against an under-estimated ψ. Proof: the homogeneous gap dynamics have the single root $r = 1 - \psi\lambda$; $|r| < 1 \iff 0 < \psi\lambda < 2$, and $0 \le r < 1 \iff 0 < \psi\lambda \le 1$ (monotone). The stationary variance follows as $\sigma^2 \cdot \Sigma_{k\ge0} r^{2k} = \sigma^2/(1 - r^2)$, minimized at $r = 0$ ($\psi\lambda = 1$). $\psi = B/M^T$ follows from $P_t = M^T_t \cdot V/Y_t$: an injection of $\lambda x_{t-1}B$ dollars raises M^T, and hence P, by $\lambda x_{t-1}B/M^T$ to first order. ∎

Proposition 4 (bounded labor-supply response and growth robustness). A worker with Frisch elasticity ν chooses hours h to maximize $\ln c - \chi \cdot h^{1+1/\nu}/(1+1/\nu)$ subject to $c = \omega h + Y_F$, where Y_F is spendable non-labor income from the floor system and $h\star$ solves the problem at $Y_F = 0$. Let $b \equiv Y_F/(\omega h\star)$ and $\ell \equiv h/h\star$. Then (i) ℓ solves $\ell^{1+1/\nu} + b \cdot \ell^{1/\nu} = 1$, strictly decreasing in b, with $d\ell/db|_{b=0} = -\nu/(1+\nu)$; (ii) a liquid floor gives $b = b_{liq} \approx rF/(\omega h\star)$, whereas locking the principal and distributing only $Y_F = \rho_{eff} \cdot F$ gives $b = (\rho_{eff}/r) \cdot b_{liq}$, with $\rho_{eff} \to 0$ as the distributed share $\kappa_d \to 0$ under full bequest; (iii) b is invariant to the level of the balanced-growth path, so the floor is a one-time level effect on labor and output, and trend growth changes by at most $\varepsilon_g \cdot [\nu/(1+\nu)] \cdot b$ in an endogenous-growth extension with growth–labor elasticity ε_g; (iv) when Y_F is growth-indexed, the labor–growth map $g \star g\square[1 + \varepsilon_g(\ell(b(g)) - 1)]$ has loop gain $\varepsilon_g \cdot [\nu/(1+\nu)] \cdot (d \ln b/d \ln g) < 1$ for relevant parameters, so its fixed point is locally stable and self-correcting. Proof in A.7.

Proposition 5 (the slow lever cannot stabilize through feedback). Augment the gap recursion of Section 3.6 with delayed feedback from a slow lever of gain λ_L and delay $d \ge 1$: $x_t = (1-\psi\lambda)x_{t-1} - \lambda_L x_{t-d} + \varepsilon_t$, with characteristic polynomial $P(z) = z^d - (1-\psi\lambda)z^{d-1} + \lambda_L$. Then (i) the small-gain box $|1-\psi\lambda| + |\lambda_L| < 1$ is sufficient for stability at every d; (ii) the critical gain $\lambda_L^{crit}(d)$ above which P has a root outside the unit circle is non-increasing in d and converges to $1 - |1-\psi\lambda|$ from above; and (iii) at any fixed $\lambda_L \in (0,1)$ in the stable region the dominant root satisfies $\max|z| \to 1$ as $d \to \infty$, with damping $1 - \max|z| = |\ln \lambda_L|/d + o(1/d)$. Hence a lever of decade-scale delay contributes only vanishingly damped modes and cannot provide effective feedback stabilization; treating L_t as a predetermined (feedforward) schedule, as Section 3.6 does, is the stability-preserving choice. Proof in A.8.

Proposition 6 (linearized dynamics and the stabilization properties). Let deviations from the balanced path follow $s_t = A s_{t-1} + e_t$ with $s_t = (x_t, \tilde{y}_t, v_t)'$ and $x_t = (1-\psi\lambda)x_{t-1} + \kappa \tilde{y}_{t-1} + \lambda_{leak} v_{t-1} + e^{\pi}_t$, $\tilde{y}_t = \varphi_y \tilde{y}_{t-1} - \gamma x_{t-1} + (1-\omega_F)e^{y}_t$, $v_t = \varphi_v v_{t-1} + e^{v}_t$, where ω_F is the acyclical floor-income share and λ_{leak} the asset → consumer leak. Then (i) the system is stable iff $|\varphi_v| < 1$ and the $(x, \tilde{y})$ block has both eigenvalues in the unit disk, and it reduces to the Proposition 3′ recursion when $\kappa = \gamma = \lambda_{leak} = 0$; (ii) the output-gap response to a demand shock is proportional to $(1-\omega_F)$, so the acyclical floor cushions demand shocks and the trough shrinks linearly in the floor-income share; (iii) the price-path gap returns to target at the KI rate $\approx 1-\psi\lambda$, while with KI off ($\psi\lambda = 0$) it carries a unit root and does not self-correct; (iv) the consumer-price response to an asset-circuit shock is bounded in steady state by $\lambda_{leak}/(\psi\lambda)$ and vanishes with the leak, so asset-circuit shocks are contained. Proof and impulse responses in A.9 (Figure 4).

Proposition 7 (forward-looking determinacy under the money-quantity anchor). Let transactional money demand take the Cagan form $m^T_t - p_t = -\alpha E_t[p_{t+1} - p_t] + y_t$ with $\alpha > 0$, and let KI set the transactional quantity in response to the price-path gap $x_t = p_t - p^*_t$ with gain $\varphi \geq 0$. Then (i) the gap follows $E_t x_{t+1} = \theta x_t + \xi_t$ with $\theta = 1 + (1+\varphi)/\alpha$; since x_t is non-predetermined, the rational-expectations equilibrium is determinate iff $|\theta| > 1$, which holds for every $\alpha > 0$ and $\varphi \geq 0$ — the price level is determinate without a Taylor-type condition. (ii) Under an interest-rate instrument responding to the gap with the same gain, the root is instead φ, so determinacy requires the Taylor principle $\varphi > 1$. (iii) Adjoining the asset circuit as a second forward-looking variable with present-value valuation, the two-jump system is determinate (two explosive eigenvalues) for asset↔consumer coupling below a threshold (≈ 0.13 at the baseline), inside which the calibrated leak (≈ 0.03) lies. Proof and Figure 5 in A.10.

Proposition 8 (welfare-optimal dividend share). A citizen with log utility and Frisch elasticity ν allocates each period's growth issuance G between a current dividend $\kappa_d G$ — consumed in the work phase and, because the floor is locked, the only component that distorts labor ($b = \kappa_d G$, Proposition 4) — and the floor, retained at $(1-\kappa_d)G$, compounded at r, and valued in the terminal phase with effective weight Δ (one if self-consumed, $\gamma_b < 1$ if bequeathed and weakly valued). Let $R \equiv \beta\Delta(1+r)$ be the retention premium, β the time-preference factor. Then (i) accumulated wealth is maximized at $\kappa_d = 0$; (ii) welfare is maximized at $\kappa_d^* = \max\{0, \min\{1, [1 + G(1+r) - R\cdot\ell(\kappa_d^* G)]/[G((1+r)+R)]\}\}$, interior for any conventional positive discounting; (iii) κ_d^* is decreasing in R — hence decreasing in patience β and in the return r, and increasing as the floor shifts from self-consumption toward weakly-valued bequest (decreasing in Δ) — and the wealth-maximizing corner $\kappa_d^* = 0$ obtains only for $R \gtrsim 1.2$, i.e. effectively negative net discounting; (iv) by the envelope theorem the Proposition 4 labor distortion is second-order at the private labor optimum, entering κ_d^* only through the budget, and becomes first-order (lowering κ_d^*) only under a labor wedge. Proof and Figure 6 in A.11.

Proposition 9 (robustness of circuit separation to bank credit). Let banks extend credit against pledgeable asset wealth at loan-to-value m, with a fraction φ_{liq} of asset wealth liquid (the locked floors being non-pledgeable) and a share χ_c of new credit spent into the transactional circuit. Then (i) bank credit adds an asset→consumer coupling $\chi_c\cdot m\cdot\varphi_{liq}$ to the structural leak, for a total coupling $\lambda_{leak} + \chi_c\cdot m\cdot\varphi_{liq}$; (ii) circuit separation — equivalently the determinacy of Proposition 7 — holds iff this total stays below the coupling threshold $\zeta^* \approx 0.13$, i.e. iff the credit intensity $\kappa_{bank} \equiv m\cdot\varphi_{liq} < (\zeta^* - \lambda_{leak})/\chi_c \approx 0.32$ under the conservative (symmetric, financial-accelerator) reading; (iii) because the floors are locked and non-pledgeable, φ_{liq} is small (≈ 0.15 at the baseline), so the critical loan-to-value $m^* = \kappa_{bank}^*/\varphi_{liq}$ exceeds 2 — no feasible leverage breaks separation — whereas fully pledgeable wealth ($\varphi_{liq} = 1$) would break it at an ordinary $m \approx 0.32$; (iv) the rule-bound structural buyer damps the accelerator, raising the critical intensity to ≈ 1.7, an independent second safeguard. Proof and Figure 7 in A.12.

A.3 Proof sketch — the rebalancing channel (A2)

A seller with wealth $W = m + e$ and target share $\mu^* = m/W$ consumes $c = \kappa_c\cdot W$. An FDCA purchase of Δ sets money to $m + \Delta$, equity to $e - \Delta$, leaving W unchanged; hence $\Delta c = \kappa_c\cdot\Delta W = 0$ (zero direct MPC). The post-sale money share $(m+\Delta)/W > \mu^*$, so the seller rebalances Δ back into equity. Aggregating, Δ remains in M^A; for aggregate holders to retain target share μ^*

of a larger asset stock, asset values rise by Δ/μ^* (multiplier $1/\mu^*$). The only leak into goods prices is the asset-wealth effect on unconstrained holders:
$\text{Leak}_t = \kappa_W \cdot (\Delta_t/\mu^*) \cdot (1 - s_t)$,
which → 0 as $s_t \to 1$ (appreciation increasingly trapped in non-spending locked floors). Thus (A2) holds exactly in the mature limit and with a bounded, quantified deviation otherwise. ∎ (sketch)

A.4 Proof sketch — the no-net-saving condition (A3 / Result 2)

A life-cycle worker targets total equity exposure E^*. Retirement equity is the locked floor F_T plus accumulated discretionary saving $\Sigma\, b_t$. Optimal $b_t > 0$ only if $F_T < E^*$. Under mature calibration F_T is approximately $413K in Mode B on the realizable basis (launch-year purchasing power; Neo-Solon 2026a/2026b, §6.7), so $F_T \geq E^*$ for the median worker; hence $b_t = 0$, labor income is consumed in M^T, and the net $M^T \to M^A$ flow is zero. The condition fails for (i) immature-system cohorts with $F_T < E^*$ (the same small-s_t window as A.3) and (ii) high-target workers with $E^* > F_T$, whose flow is a bounded level shift, not regime-induced. ∎ (sketch)

A.5 Calibration and falsification

The model is analytical; its results are conditional on five quantities left for empirical work. Each is a falsification point — a value or path that, if it came out wrong, would break the corresponding result.

1. **Price-leak magnitude (κ_W, μ^*, s_t path).** Calibrate κ_W from the wealth-effect literature (central ≈ 0.03, low for the asset-rich sellers who execute the rebalancing; see 3.2 calibration note), μ^* from household money/asset shares, and s_t from the projected locked-float trajectory. The model predicts the leak $\kappa_W \cdot (\Delta/\mu^*) \cdot (1 - s_t)$ is small at launch and declining; an observed leak that did *not* decline as floors matured would falsify the maturation mechanism.
2. **Return-compression premium (A^*, φ).** Calibrate steady-state net absorption A^* (gross K1/K2 purchases net of drawdown L_t at demographic maturity) and equity-supply elasticity φ from issuance-response data. The model predicts a finite premium A^*/φ and $A^* \to$ small; a permanently large A^* with $\varphi \approx 0$ would falsify Proposition 2.
3. **Circuit-tracking (M^T vs Y).** Establish that the issuance rule plus liquidation flow can hold M^T growth to real output growth across the cycle. Sustained divergence of M^T from Y not attributable to a measured shock would falsify the Proposition 1 stability condition.
4. **Dividend-rule pass-through (ψ).** The path-targeting stability condition of Proposition 3 is $\psi\lambda < 1$, where ψ is the within-period pass-through from path-closure issuance to realized inflation and λ the chosen closure gain. The model predicts $\psi\lambda < 1$ (monotone closure). A measured ψ such that $\psi\lambda \geq 1$ at the baseline $\lambda = 0.5$ — i.e. the dividend rule observed to overshoot and oscillate around the price path rather than close it monotonically — would falsify the self-correction claim. Because λ is a chosen gain, this is a correctable falsification: the rule re-tunes to $\lambda < 1/\psi$ to restore stability.

Pass-through structure ($\psi = M2/M^T$). Proposition 3′ identifies the stability-critical ratio as $\psi = B/M^T$, where B is the base of the gap-closure injection. Under the transactional indexing adopted in Section 4.4 ($B = M^T$), $\psi \equiv 1$ and the condition is the maturity-invariant $\lambda < 1$. A transactional share $M^T/M2$ that fell over time while the rule was (mis)indexed to M2 would drive $\psi\lambda$ through the oscillation and divergence thresholds; the falsification is correctable by re-basing the injection on M^T or adopting the adaptive gain $\lambda_t = c \cdot M^T_t/M2_t$.

5. **Money-share return sensitivity (μ_r).** The §3.2 endogenous-μ^* correction predicts the desired money share rises as equity returns fall, i.e. $\mu_r \equiv -\partial\mu^*/\partial r > 0$. This is a *sign*

restriction, not a point value: the model relies only on $\mu_r > 0$ to conclude that endogenizing μ^* shrinks rather than enlarges the price leak. A measured $\mu_r \leq 0$ (money demand flat or rising in returns) would remove this favorable second-order correction. Because the headline leak bound of §3.2a deliberately omits the correction and uses the larger fixed-μ^* value, a wrong-signed μ_r does not loosen the headline result; it only forfeits a conservatism margin.

6. **Labor-supply response (ν, b).** ***The wealth effect on labor is governed by the Frisch elasticity ν and the spendable floor-income ratio b — not by the floor stock (Proposition 4). Calibrate ν from the labor-supply literature (central ≈ 0.5) and b from the distribution design (κ_d, drawdown schedule, bequest fraction). A measured labor-supply decline materially exceeding $[\nu/(1+\nu)]\cdot b$, or a b that rises with the floor stock rather than staying scale-invariant, would falsify the level-effect reading and reopen the growth-collapse objection.***
7. **Dynamic responses (κ, γ, ω_F, λ_{leak}).** ***The impulse responses of Proposition 6 are governed by the output→price slope κ, the price-gap→output cooling γ, the floor-income share ω_F, the asset persistence φ_v, and the leak λ_{leak}. A demand-shock output trough that fails to shrink with the floor-income share, a cost-push gap that persists under an active KI, or an asset-circuit shock that passes more than $\lambda_{leak}/(\psi\lambda)$ into consumer prices would each falsify a corresponding claim.***
8. **Forward-looking determinacy (α, coupling).** ***Determinacy of the price path turns on the money-demand semi-elasticity α and the asset↔consumer coupling (Proposition 7). Any $\alpha > 0$ secures determinacy under the money-quantity rule, but a coupling exceeding ≈ 0.13 would tip the system into indeterminacy; estimating the leak below that threshold is the falsifiable content.***
9. **Bank-credit coupling (m, φ_{liq}, χ_c).** ***Circuit separation under endogenous credit (Proposition 9) is falsifiable through the credit intensity $m\cdot\varphi_{liq}$ and spend-through χ_c: separation requires $\lambda_{leak} + \chi_c\cdot m\cdot\varphi_{liq}$ below the coupling threshold $\zeta^* \approx 0.13$. The pledgeable fraction φ_{liq} — bounded well below one by the non-pledgeability of locked floors — and the loan-to-value m are the estimable quantities.***
10. **Dividend optimum (β, Δ, r).** ***The welfare-optimal dividend share turns on the time-preference factor β, the terminal-value weight Δ (self-consumption vs bequest), and the return r through the retention premium $R = \beta\Delta(1+r)$ (Proposition 8). The wealth-maximizing default $\kappa_d = 0$ is welfare-optimal only for $R \gtrsim 1.2$; a social objective with positive net discounting paired with a $\kappa_d = 0$ policy identifies an explicit accumulation-over-consumption value choice rather than a welfare optimum.***

A replication exercise calibrating these against US data, in the spirit of the empirical paper (Neo-Solon, 2026b), is the natural next step and is not attempted here.

A.6 Demographic equity-flow model (dating Proposition 2)

Section 6.6 rests on an overlapping-generations simulation of the floors' net external equity flow. Each annual cohort accumulates a Stable Floor from K1 and K2 deposits compounding at the equity total return r; from age 65 it draws consumption at rate w; at death the residual passes to a beneficiary's floor under the inheritance rule and continues to compound. The price-pressure-relevant quantity is net external absorption — issuance into floors minus retiree consumption. Dividends are not counted as independent demand: in aggregate a reinvested dividend buys back exactly the ex-dividend price drop it is funded from, a wash; only consumed dividends, already part of retiree withdrawals, leave the circuit.

At the central calibration ($r = 4.5\%$, $w = 4\%$, retirement 65, mortality 85, $g = 1.5\%$, $n = 0.5\%$), net absorption is positive early — about +1.3% of GDP in the second decade — and crosses to negative near year 55 as the retiree stock matures, settling into a structural outflow thereafter. The crossover is robust — the buyer-to-seller transition always occurs — and varying each parameter individually across withdrawal rates of 3–5%, returns of 3.5–6.5%, population growth of 0% to 1%, and mortality ages of 80–90 places it between years 43 and 67, wider under combined extremes; only the magnitude of the mature outflow varies. That magnitude is the $r > g$ rebalancing: the excess by which floor holdings outgrow GDP, equal at a bounded market share to $(r - g)$ times the floor stock — a few percent of GDP per year — and zero when r does not exceed g. Replacing inheritance with liquidation at death adds a death-liquidation spike that roughly doubles the mature outflow and makes it lumpy; inheritance keeps it smooth and bounded.

The result is therefore the dated form of Proposition 2 — a net-buyer accumulation phase followed by the structural-seller steady state — not a contradiction of it, and the mature outflow is the net demand the Q-channel absorbs through firm issuance (Section 6.2). It does not support a stronger claim that the floors are permanent net buyers. The model and its robustness battery are reproducible from demographic_flow_model.py in the empirical replication package, which also generates the net-flow figure.

A.7 Proof sketch — labor supply and growth robustness

Write per-period utility $\ln c - \chi \cdot h^{(1+1/\nu)}/(1+1/\nu)$ with budget $c = \omega h + Y_F$. The intratemporal first-order condition is $\omega/c = \chi \cdot h^{(1/\nu)}$. At the benchmark $Y_F = 0$ this fixes χ, and substituting with $\ell = h/h_\star$, $b = Y_F/(\omega h_\star)$ reduces the condition to $\ell^{(1+1/\nu)} + b \cdot \ell^{(1/\nu)} = 1$. The left side is increasing in ℓ and in b, so ℓ is strictly decreasing in b; implicit differentiation at $b = 0$, $\ell = 1$ gives $(1+1/\nu)\ell' + 1 = 0$, i.e. $d\ell/db = -\nu/(1+\nu)$. This is part (i).

For (ii), a liquid principal contributes its full annuitized return to spendable income, $Y_F = rF$, so $b_{liq} = rF/(\omega h_\star)$; a locked principal contributes only the distributed flow $Y_F = \rho_{eff} \cdot F$, whence $b = (\rho_{eff}/r) \cdot b_{liq}$. The effective yield ρ_{eff} collects the growth-indexed distribution (scaling with κ_d) and the bequest-net annuitized retirement value; as $\kappa_d \to 0$ with the inheritance rule diverting the residual to the next floor, $\rho_{eff} \to 0$ and the labor effect vanishes.

For (iii), along a balanced path ω and F both grow at g (the floor is growth-matched), so $b = Y_F/(\omega h_\star)$ is independent of the path's level. A change in b therefore shifts the level of labor and output without altering trend growth; embedding $g = g_\star[1 + \varepsilon_g(\ell - 1)]$ for an endogenous-growth elasticity ε_g bounds the trend change by $\varepsilon_g \cdot [\nu/(1+\nu)] \cdot b$.

For (iv), let the distributed flow be growth-indexed, $b = b(g)$ with $d \ln b/d \ln g = 1$, and compose the map $g \mapsto g_\star[1 + \varepsilon_g(\ell(b(g)) - 1)]$. Its derivative at the fixed point has magnitude $\varepsilon_g \cdot [\nu/(1+\nu)] \cdot (d \ln b/d \ln g) \cdot b$, below one for empirically relevant ε_g and b (≈ 0.016 at $\nu = 0.5$, $\varepsilon_g = 0.5$, $b = 0.1$), so the fixed point is a locally stable attractor and the system self-corrects after a growth disturbance. With ω equal to the marginal product of labor, a fall in h raises ω and lowers b, reinforcing the contraction. ∎

A.8 Proof sketch — delayed feedback and the two-speed design

Setting $\varepsilon = 0$ and $x_t = z^t$ in $x_t = (1-\psi\lambda)x_{t-1} - \lambda_L x_{t-d}$ gives, after dividing by z^{t-d}, the characteristic polynomial $P(z) = z^d - (1-\psi\lambda)z^{d-1} + \lambda_L$. At $\lambda_L = 0$ it factors as z^{d-1}

(z − (1−ψλ)): one pole at the fast-loop value 1−ψλ and d−1 poles at the origin — the first-order recursion of Section 3.6, stable iff 0 < ψλ < 2.

For (i), iterating the recursion gives |x_t| ≤ |1−ψλ|·sup|x_(t−1)| + |λ_L|·sup|x_(t−d)|, so if |1−ψλ| + |λ_L| < 1 the map is a contraction in the sup norm and x_t → 0 regardless of d. The stable region therefore contains the small-gain box at every delay.

For (ii), the d−1 origin poles migrate outward as λ_L rises; the product of the roots of P equals (−1)^d λ_L, so their geometric-mean modulus is |λ_L|^(1/d). The first boundary crossing defines λ_L^crit(d), and a numerical root scan (supplement, verify_proposition_5.py) gives it decreasing monotonically — 1.00 at d ≤ 2, 0.93 at d = 4, 0.91 at d = 8 — toward the small-gain value 1 − |1−ψλ| (here 0.90) as d → ∞. The extra margin a slow feedback lever could exploit thus vanishes with its delay.

For (iii), |λ_L|^(1/d) → 1 for any fixed λ_L ∈ (0,1), so the dominant root approaches the unit circle: max|z| = exp(−|ln λ_L|/d)(1 + o(1)), i.e. damping 1 − max|z| = |ln λ_L|/d + o(1/d). At decade-scale delay (d on the order of 10–40 steps) the loop, though stable, decays only over many delay periods and rings at the delay frequency, as Figure 3(c) shows. A feedforward schedule adds no such pole and preserves the fast loop's clean geometric decay. ∎

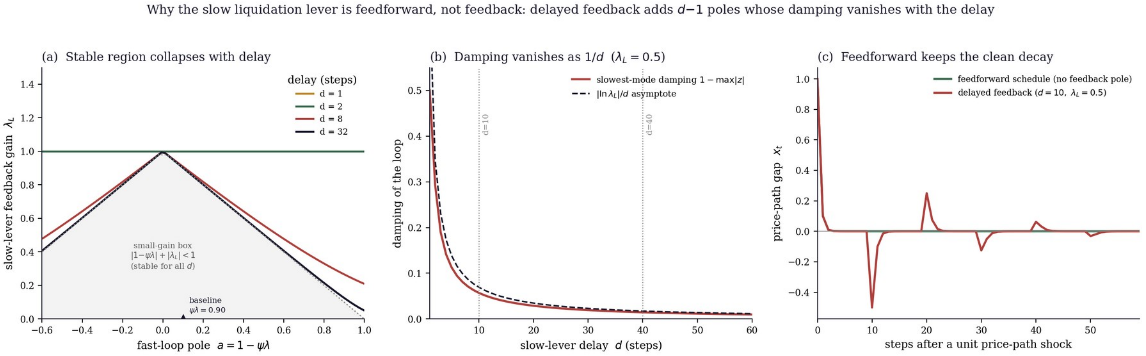

Figure 3. Why the slow liquidation lever is feedforward, not feedback. (a) the stability region in the (a = 1−ψλ, λ_L) plane collapses toward the delay-independent small-gain box |1−ψλ| + |λ_L| < 1 as the delay d grows; (b) at a fixed feedback gain (λ_L = 0.5) the damping of the slowest mode vanishes like |ln λ_L|/d, so a decade-scale delay (d ≈ 10–40) leaves almost none; (c) impulse responses to a unit price-path shock — a feedforward schedule keeps the fast loop's clean decay, while delayed feedback rings at the delay period and decays only slowly. Baseline ψλ = 0.90.

A.9 Linearized dynamics and impulse responses (Proposition 6)

Write the system as s_t = A s_(t−1) + e_t with A = [[1−ψλ, κ, λ_leak], [−γ, φ_y, 0], [0, 0, φ_v]]. The matrix is block triangular: the asset row decouples with eigenvalue φ_v, and the 2×2 block B = [[1−ψλ, κ], [−γ, φ_y]] governs (x, ỹ). Stability requires |φ_v| < 1 and both eigenvalues of B inside the unit circle (|tr B| < 1 + det B < 2). At the baseline (ψλ = 0.9, κ = 0.08, γ = 0.15, φ_y = 0.7) the eigenvalues are 0.12, 0.68, 0.90; setting κ = γ = λ_leak = 0 returns x_t = (1−ψλ)x_(t−1) + e^π_t — Proposition 3′. This is (i).

For (ii), a demand disturbance e^y enters the output row with loading (1−ω_F), since the floor-income share ω_F is acyclical and does not transmit the disturbance. The response of ỹ is the B-propagated path of (1−ω_F)e^y, linear in (1−ω_F); the trough scales as (1−ω_F), from −1.00 at ω_F = 0 to −0.85 at ω_F = 0.15. With κ small and KI active the induced price-path gap stays near zero.

For (iii), a cost-push shock sets x_0 = 1; the homogeneous return is governed by the dominant eigenvalue of B, ≈ 1−ψλ for small κγ, so x_t ≈ (1−ψλ)^t → 0. With

KI disabled (ψλ = 0) the (1,1) entry becomes one and B acquires a unit root; the gap then fails to return (x_20 ≈ 0.46 in simulation versus ≈ 0 with KI).

For (iv), an asset shock gives v_t = φ_v^t and feeds x only through λ_leak v_(t−1), so x_t = (1−ψλ)x_(t−1) + λ_leak v_(t−1). For a sustained unit asset deviation the steady-state consumer-price gap is λ_leak/(ψλ), which also bounds the peak response — about 0.03 at the baseline. Asset-circuit shocks are contained: the dynamic counterpart of the circuit separation of Section 3. ∎

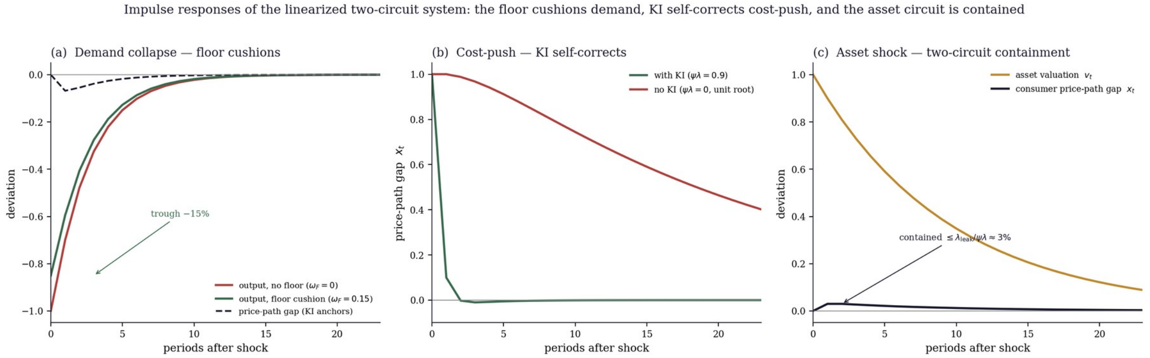

Figure 4. Impulse responses of the linearized two-circuit system. (a) a negative demand shock — the acyclical floor cushions the output trough (15% shallower at ω_F = 0.15) while KI holds the price-path gap near zero; (b) a cost-push shock — with KI active the gap returns at rate ≈ 1−ψλ, while with KI off it carries a unit root and persists; (c) an asset-circuit shock — valuations move fully but the consumer price-path gap is contained to ≤ λ_leak/(ψλ) ≈ 3%. Baseline ψλ = 0.9, κ = 0.08, γ = 0.15, φ_y = 0.7, φ_v = 0.9, λ_leak = 0.03.

A.10 Forward-looking determinacy (Proposition 7)

Subtract the target-path identity from the Cagan demand and the rule to write both in gaps: money demand μ_t = (1+α)x_t − α E_t x_(t+1) + ṽ_t, and the KI quantity rule μ_t = −φ x_t + u_t. Equating and solving for the lead gives E_t x_(t+1) = [(1+α+φ)/α] x_t + (ṽ_t − u_t)/α, i.e. θ = (1+α+φ)/α = 1 + (1+φ)/α.

Since the price level adjusts freely, x_t is non-predetermined; Blanchard and Kahn (1980) require, for a unique non-explosive solution, exactly one eigenvalue outside the unit circle — here |θ| > 1. As (1+φ)/α > 0 for all admissible α > 0, φ ≥ 0, the condition holds universally, and the unique bounded solution is the forward sum x_t = −Σ_(j≥0) θ^(−(j+1)) E_t ξ_(t+j). This is (i).

For (ii), a cashless interest-rate instrument i_t = φ x_t with the Fisher relation i_t = E_t x_(t+1) (constant real rate, in deviations) gives E_t x_(t+1) = φ x_t, whose root is φ; determinacy requires |φ| > 1, the Taylor principle. The difference is structural: the money-quantity instrument pins the price level directly, supplying the 1 + 1/α margin the interest-rate instrument lacks.

For (iii), adjoin an asset-valuation gap q_t with q_t = a E_t q_(t+1) + b x_t, a = 1/(1+r) ∈ (0,1), giving E_t[x_(t+1), q_(t+1)]′ = M [x_t, q_t]′ with M = [[θ, −ζ], [−b/a, 1/a]]. Two non-predetermined variables require two eigenvalues outside the unit circle. At zero coupling these are θ > 1 and 1/a > 1; by continuity both remain outside for small coupling, and a numerical scan (supplement, verify_proposition_7.py) places the boundary at ζ = b ≈ 0.13. The calibrated leak (≈ 0.03) lies inside, so the two-circuit equilibrium is determinate — a further dividend of circuit separation. ∎

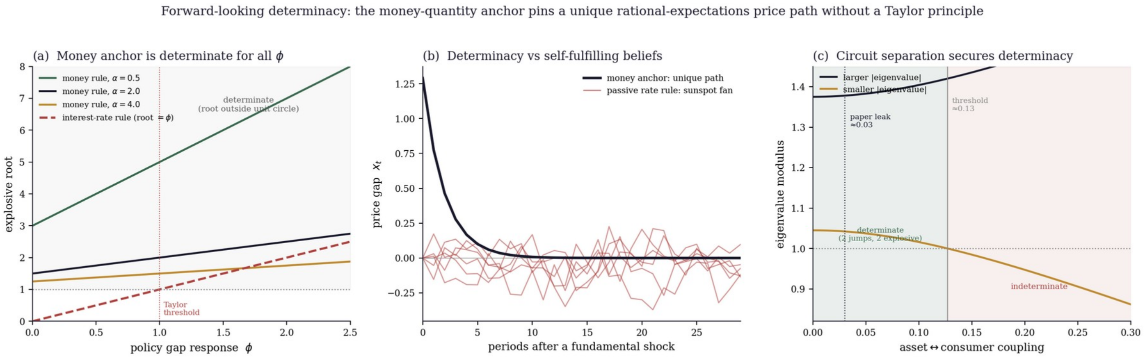

Figure 5. Forward-looking determinacy under the money-quantity anchor. (a) the explosive root against the gap response φ: under the money-quantity rule $\theta = 1 + (1+\varphi)/\alpha$ exceeds one for every α and φ (determinate), while the interest-rate analog has root φ and needs the Taylor principle $\varphi > 1$; (b) a fundamental shock yields a unique bounded price path under the money anchor, whereas a passive interest-rate rule admits a fan of self-fulfilling (sunspot) equilibria; (c) with the asset circuit adjoined, the two-jump system stays determinate (two explosive eigenvalues) until the asset ↔ consumer coupling reaches ≈ 0.13, inside which the calibrated leak (≈ 0.03) lies. Baseline $\alpha = 4$, $\varphi = 0.5$, $r = 4.5\%$.

A.11 The welfare-optimal dividend share (Proposition 8)

Write lifetime welfare $W(\kappa_d) = \ln c1 - \chi h^{(1+1/\nu)}/(1+1/\nu) + \beta\Delta \ln c2$ with $c1 = h + \kappa_d G$ (work-phase labor income plus dividend) and $c2 = 1 + (1-\kappa_d)G(1+r)$ (retirement baseline plus the compounded retained floor), and $b = \kappa_d G$. Because the floor is locked, only the current dividend enters the labor wealth effect, so $h = \ell(\kappa_d G)$ solves the Proposition 4 condition. Accumulated wealth $(1-\kappa_d)G(1+r)$ is maximized at $\kappa_d = 0$ — claim (i).

Differentiating and using the intratemporal labor optimum $u'(c1) = \chi h^{(1/\nu)}$ — which cancels the $dh/d\kappa_d$ terms (the envelope theorem, claim (iv)) — the first-order condition is $(1/c1)G = \beta\Delta(1/c2)G(1+r)$, i.e. $c2/c1 = R$ with $R \equiv \beta\Delta(1+r)$. Substituting the budget identities yields $\kappa_d^* = [1 + G(1+r) - R\,\ell(\kappa_d^* G)]/[G((1+r)+R)]$, clamped to $[0,1]$ — claim (ii); the optimum is interior whenever R lies below the corner value.

For (iii), $\partial\kappa_d^*/\partial R < 0$ by inspection of the first-order condition; since $R = \beta\Delta(1+r)$, κ_d^* falls with patience β, with the return r, and with the terminal-value weight Δ, and rises with their opposites. A numerical scan (supplement, verify_proposition_8.py) places the corner $\kappa_d^* = 0$ at $R \approx 1.21$ for the baseline G and ν; as $R = \beta\Delta(1+r) \leq 1$ for any $\beta \leq 1$, $\Delta \leq 1$ at $r = 4.5\%$, the corner requires effectively negative net discounting.

For (iv), the labor distortion enters W only through c1 in the budget; at the private labor optimum its marginal welfare contribution is second-order, so κ_d^* is governed by the intertemporal–bequest tradeoff rather than the distortion. A wedge between the private and social value of labor (a tax τ or a production externality) adds a first-order term that raises the marginal cost of distribution, lowering κ_d^*. ∎

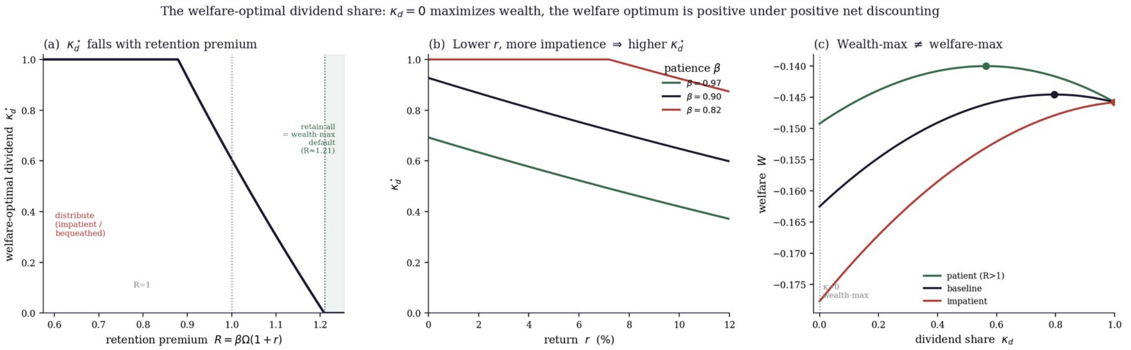

Figure 6. The welfare-optimal dividend share. (a) the optimal share κ_d^* against the retention premium $R = \beta\Delta(1+r)$: κ_d^* falls as retention becomes more attractive and reaches the wealth-maximizing corner $\kappa_d^* = 0$ only near $R \approx 1.2$; (b) κ_d^* against the return r for three patience levels β, falling with r and rising with impatience; (c) welfare $W(\kappa_d)$ for patient, baseline, and impatient citizens — the optima (dots) lie strictly above the wealth-maximizing $\kappa_d = 0$ except in the patient limit. Baseline $G = 0.20$, $\nu = 0.5$, $\Delta = 1$, $r = 4.5\%$.

A.12 Robustness of circuit separation to bank credit (Proposition 9)

Model bank credit in reduced form. Banks lend against pledgeable asset wealth: with loan-to-value m and a liquid (pledgeable) fraction φ_{liq} of asset wealth, credit capacity is $L = m \cdot \varphi_{liq} \cdot A$, where A is asset-circuit wealth and the locked floors are excluded from A by construction. A share χ_c of new credit is spent on goods, entering the transactional circuit, so a unit change in asset wealth induces a transactional inflow $\chi_c \cdot m \cdot \varphi_{liq}$. Adding the structural leak λ_{leak} gives a total asset→consumer coupling $\zeta_{eff} = \lambda_{leak} + \chi_c \cdot m \cdot \varphi_{liq}$. This is (i), with $\kappa_{bank} \equiv m \cdot \varphi_{liq}$ the credit intensity.

For (ii), insert ζ_{eff} into the forward-looking system of Proposition 7, $M = [[\theta, -\zeta_{eff}], [-b_{eff}/a, 1/a]]$. Under the conservative reading credit and collateral reinforce each other (a financial accelerator), so the feedback is symmetric, $b_{eff} = \zeta_{eff}$, and determinacy fails when ζ_{eff} reaches the threshold $\zeta^* \approx 0.127$ of A.10. Hence separation requires $\chi_c \cdot \kappa_{bank} < \zeta^* - \lambda_{leak}$, i.e. $\kappa_{bank} < (\zeta^* - \lambda_{leak})/\chi_c \approx 0.32$ at $\chi_c = 0.30$, $\lambda_{leak} = 0.03$.

For (iii), the critical loan-to-value is $m^* = \kappa_{bank}^*/\varphi_{liq}$. The lock fixes φ_{liq} small: with citizen wealth concentrated in non-pledgeable floors, $\varphi_{liq} \approx 0.15$, giving $m^* \approx 2.16$ — above any feasible loan-to-value, so credit cannot break separation. Removing the lock ($\varphi_{liq} = 1$) gives $m^* \approx 0.32$, within ordinary lending. Non-pledgeability of the floor is the binding safeguard.

For (iv), the conservative reading overstates the accelerator: asset demand under the Citizens Standard is set by the rule-bound K1/K2 buyer, not by credit, so credit-fueled spending does not raise valuations one-for-one and the consumer→asset feedback b_{eff} stays near its structural value λ_{leak} rather than tracking ζ_{eff}. Holding b_{eff} at λ_{leak}, the determinant condition pushes the critical intensity to $\kappa_{bank} \approx 1.7$, so separation survives even fully pledgeable, highly leveraged wealth. The lock and the structural buyer are independent; the baseline $\kappa_{bank} \approx 0.075$ lies inside both. A full intermediary-sector model is deferred to companion work. ∎

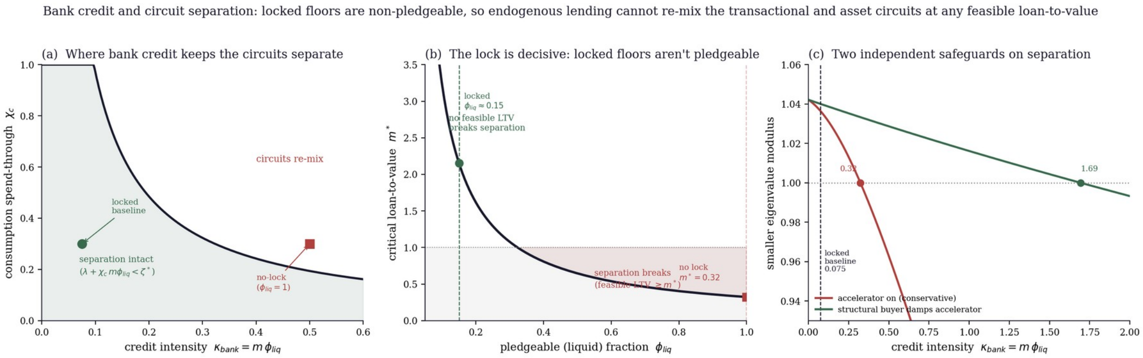

Figure 7. Bank credit and the robustness of circuit separation. (a) the region in which bank credit leaves the circuits separate, in the plane of credit intensity κ_bank = m·φ_liq and consumption spend-through χ_c: separation holds while λ_leak + χ_c·κ_bank stays below the threshold ζ ≈ 0.13; the locked baseline lies inside, the fully pledgeable (no-lock) point outside; (b) the critical loan-to-value m* = κ_bank*/φ_liq against the pledgeable fraction φ_liq — with locked floors (φ_liq ≈ 0.15) no feasible loan-to-value breaks separation, whereas fully pledgeable wealth breaks it at m ≈ 0.32; (c) the smaller eigenvalue against credit intensity under the conservative (financial-accelerator) reading and under the structural-buyer reading that damps the accelerator, giving two independent thresholds (≈ 0.32 and ≈ 1.7) with the baseline (0.075) inside both. Baseline λ_leak = 0.03, χ_c = 0.30.*

References

Blanchard, O. J., & Kahn, C. M. (1980). The Solution of Linear Difference Models under Rational Expectations. *Econometrica*, 48(5), 1305–1311.

Brayton, F., & Tinsley, P. (1996). *A Guide to FRB/US: A Macroeconomic Model of the United States.* Finance and Economics Discussion Series 1996-42, Board of Governors of the Federal Reserve System.

Bauducco, S., & Caputo, R. (2020). Wicksell versus Taylor: A Quest for Determinacy and the (Ir)relevance of the Taylor Principle. Journal of Economic Dynamics and Control, 114, 103865.

Cagan, P. (1956). The Monetary Dynamics of Hyperinflation. In M. Friedman (Ed.), *Studies in the Quantity Theory of Money* (pp. 25–117). University of Chicago Press.

Chodorow-Reich, G., Nenov, P. T., & Simsek, A. (2021). Stock Market Wealth and the Real Economy: A Local Labor Market Approach. *American Economic Review*, 111(5), 1613–1657. (NBER Working Paper No. 25959, 2019.)

Di Maggio, M., Kermani, A., & Majlesi, K. (2020). Stock Market Returns and Consumption. *The Journal of Finance*, 75(6), 3175–3219. (NBER Working Paper No. 24262, 2018.)

Dynan, K. E., & Maki, D. M. (2001). *Does Stock Market Wealth Matter for Consumption?* Finance and Economics Discussion Series 2001-23, Board of Governors of the Federal Reserve System.

Giannoni, M. P. (2014). Optimal Interest-Rate Rules and Inflation Stabilization versus Price-Level Stabilization. Journal of Economic Dynamics and Control, 41, 110–129.

Ireland, P. N. (2009). On the Welfare Cost of Inflation and the Recent Behavior of Money Demand. *American Economic Review*, 99(3), 1040–1052.

Lucas, R. E. (2000). Inflation and Welfare. *Econometrica*, 68(2), 247–274.

Piazzesi, M., Rogers, C., & Schneider, M. (2019). Money and Banking in a New Keynesian Model. Working Paper, Stanford University.

Poterba, J. M. (2000). Stock Market Wealth and Consumption. *Journal of Economic Perspectives*, 14(2), 99–118.

Sidrauski, M. (1967). Rational Choice and Patterns of Growth in a Monetary Economy. *American Economic Review*, 57(2), 534–544.

Woodford, M. (1994). Monetary Policy and Price Level Determinacy in a Cash-in-Advance Economy. Economic Theory, 4(3), 345–380.

Woodford, M. (1995). Price-Level Determinacy Without Control of a Monetary Aggregate. Carnegie-Rochester Conference Series on Public Policy, 43, 1–46.

Neo-Solon (2026a). *The Citizens Standard: One Model, Many Systems — A Constitutional Monetary Architecture.* SSRN Working Paper 6702518.

Neo-Solon (2026b). *The Citizens Standard: A Historical Counterfactual — Empirical Analysis of an Alternative US Monetary Architecture, 1960–2055.* SSRN Working Paper 6735078.

Neo-Solon (2026c). *The Citizens Standard: Transition Architecture and Migration Mechanics.* SSRN Working Paper 6810741.

Neo-Solon (2026d). *The Citizens Standard: A Statutory Implementation Pathway.*Neo-Solon (2026i). The Citizens Standard: Empirical Validation of the Transactional Aggregate — Identifying M^T Independently of the Price-Stabilization Claim. SSRN Working Paper 6973261. SSRN Working Paper 6873798.