THE CITIZENS STANDARD

External Interoperability and the Common Anchor

Computed-Rate Settlement, the Common Zero Anchor, and Domestic Feasibility under Heterogeneous Sovereign-Money Economies

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: A Macroeconomic Model of a Two-Circuit Monetary System (Neo-Solon, 2026e) · SSRN 6939418

The Citizens Standard: Full-Reserve Banking and the Two-Circuit System (Neo-Solon, 2026f) · SSRN 6939498

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

Abstract

The first six papers of the Citizens Standard (CS) specify a domestic monetary architecture. This paper closes the multilateral question: how CS economies — with different productivity levels, demographic profiles, and trade balances — interoperate when they operate side by side and must exchange goods, services, and claims. The design commitment is twofold. CS economies settle cross-border claims through a computed real-purchasing-power exchange layer: an exchange rate calculated from verifiable real data rather than traded in a market, and therefore without a speculative attack surface. (This class of settlement mechanism was introduced by Serra's EQUA; CS builds on that idea and specifies its own layer of the same class.) And CS members commit to a predictability standard: a credible, low-variance price path anchored at a common target of approximately zero inflation. The central result is that zero is the uniquely robust anchor. Heterogeneous national inflation levels distort the bilateral real rate under sticky wages; a common level reduces this, but a common positive level (such as a 2–4% corridor) leaves a residual whenever cross-country wage stickiness differs, while a common zero removes the inflation-driven distortion for every wage-adjustment process at once. Zero is, in addition, the cost-minimizing level and the non-arbitrary focal point. The domestic results are computed on the framework's actual issuance engine at its published launch calibration; the external-layer results are a calibrated property of the settlement layer's exact formula.

Methods and reproducibility. Two classes of result appear in this paper and are held to different standards. The domestic claims (Sections 5–6, 9) are computed on the CS issuance engine implemented directly from the published specification (Paper 1, §3.3 and Appendix B) at the framework's real launch parameters (M2 = $22,366B, real growth 2%, K1 = 2.5% of GDP per capita); the engine reproduces the published launch figures to the dollar. The external-layer claims (Section 7) are computed on the settlement layer's exact formula. Because that layer is a proposed mechanism with no real-world instance, those results are presented as a calibrated mechanism property — robust across the empirically observed ranges of cross-country productivity growth and wage stickiness — not as an empirical forecast. Supporting code: cs_engine.py, equa_model_v3.py, equa_redteam.py, equa_stress.py.

Supplementary replication. Because this robustness is stated for the empirically observed ranges of cross-country divergence, a companion check makes that literal: it feeds the paper's own real-rate-distortion mechanism the inflation divergences that major economies have actually produced — the 2022 spike (US and the euro area near 8–9 percent against Japan near 2.5, a gap of roughly six points), the 2021 gap, and the harder sustained case of Japan's roughly seventeen-year deflation divergence of two to three points — rather than the swept abstract shock variances of the stress battery. Under a common zero anchor the terminal bilateral real-rate distortion is approximately zero for every one of these episodes, because transient divergences around a common zero re-converge; under a common positive anchor of two percent the same divergences leave a structural wedge of about −5.8 percent. The real magnitudes are larger than the swept ones and the anchor still absorbs them. The scope of the result is that this robustness is to *transient* divergence around a held commitment: a permanent unilateral level split — a member abandoning the anchor and running k points hot forever — leaves a proportional k-point residual, so the anchor's robustness rests on the commitment being honored, which is the same governance question that attends the domestic rule, not a divergence the mechanism neutralizes on its own. Runnable as the anchor-real-shocks module in the distribution replication package.
JEL classification: F31, F33, F02, E31, E42, E52

Keywords: Citizens Standard; external interoperability; computed exchange rate; common price anchor; zero inflation; wage stickiness; sovereign money; real purchasing power

Contents

1. The Multilateral Problem

A single sovereign-money economy under CS has a natural tendency, established in the earlier papers, toward mild productivity-driven deflation: with the money supply governed rather than borrowed into existence, real growth lowers prices unless the issuer chooses otherwise. In isolation this is benign and indeed mildly optimal.

The multilateral case introduces a constraint domestic analysis cannot see. If economy A runs near-zero or mildly deflationary while its trading partner B runs higher inflation, the difference must be absorbed somewhere — in the exchange rate, in trade flows, or in internal adjustment. The question is whether heterogeneous national price paths are compatible with a stable, manipulation-resistant exchange system, and if not, what minimum coordination restores compatibility without surrendering monetary sovereignty. The paper answers in three steps: the settlement layer that prices the relationship (Section 3), the anchor members commit to (Section 4), and the domestic architecture that makes the anchor cheap and credible to hold (Sections 5–6).

2. The Interoperability Requirement

Two distinct properties of a member's price path matter, and they must not be conflated.

Variance of the differential is always binding. The exchange layer prices the differential between members' inflation paths. If that differential is predictable, the resulting exchange-rate drift is anticipated and priced; it is not destabilizing. If inflation is volatile, the drift is unpredictable and the rate inherits the volatility. What matters is the variance of the differential, which falls when shocks are common across members and rises when they are idiosyncratic (Section 7, result 6). The system therefore requires every member to deliver a low-variance, credible path. This requirement is independent of the level.

Level is binding only through wage stickiness, and only a zero level is fully robust. Under full wage indexation the level is irrelevant (Section 7, result 2). Under realistic, incomplete indexation, heterogeneous levels generate a real distortion — but a modest, bounded one (result 3). A common level reduces the distortion; however, a common positive level leaves a residual whenever members' wage stickiness differs (result 4). Only a common zero anchor removes the inflation-driven distortion across all wage-adjustment processes, because at zero inflation there is no wage–price gap for stickiness to act on. The level question therefore does not merely select a common number on cost grounds; it selects zero on robustness grounds.

3. The External Settlement Layer

CS specifies, and commits to interoperate through, an external settlement layer of the computed real-purchasing-power class. This class of mechanism was introduced by Serra's EQUA, which CS builds on. For two monetary areas A and B the layer defines a rate equal to the ratio of labor-hours required to purchase a fixed, constitutionally defined basket of essential goods and services in each economy:

$$E(A,B) = H_A / H_B$$

where H is the number of hours of labor at median wage required to purchase the basket. The rate is continuously updated by an inflation-differential component — the relative purchasing-power-parity factor $(1 + \pi A)/(1 + \pi B)$ — and a bounded real-trade-flow component that cannot override the first. The rate is published as a display, not a market: there is no spread, no position to take, and so no speculative attack surface. Three properties matter for CS:

1. **It prices heterogeneous-but-credible members.** Because the rate is computed from each member's verifiable data, members need not share a number to interoperate; they need a predictable path.
2. **Its inflation component is real-neutral.** The relative-PPP update passes the inflation differential to the nominal rate and leaves the real rate untouched (Section 7, result 1). Inflation differences never reach the real rate through the formula; any residual effect operates through member wage-setting, not through the layer.
3. **Its anchor H is a real measure.** Labor-hours for a fixed basket is invariant to inflation once wages have adjusted, and falls with genuine productivity. The layer thus delivers the correct real adjustment — a high-productivity economy's currency strengthens — without discretionary intervention.

CS does not claim originality for this class of layer; it builds on the EQUA mechanism that introduced it. CS's own contribution is the domestic architecture that supplies the layer with the low-variance credibility it requires (Sections 5–6) and the common anchor members hold (Section 4).

4. The Common Anchor: Price Stability at Zero

Each CS member commits to a price path that is (i) low-variance and credible and (ii) anchored at a common target of approximately zero inflation — the same target for every member, not a level each chooses. The operating definition is a common point target near zero, with a narrow tolerance band; a deflation buffer that lifts the operating mean slightly above zero is permissible and carries only a small, bounded cost (Section 7). The anchor is therefore predictability plus a common zero point — chosen for robustness first, cost and focality second. A member meets this anchor by operating a price-stable Mode — Mode B or Mode T, or Mode A within the deflation-buffer tolerance; Mode C, whose approximately +2% target funds a visible dividend through the KI channel, falls outside the common zero anchor. No member need adopt Mode C to pay a dividend, however, because K3 delivers an equivalent citizen dividend at zero inflation (Section 5) — which is precisely why the common-zero anchor is inexpensive for CS economies.

5. Why CS Can Hold the Anchor: the Channel Architecture

The anchor in Section 4 is a requirement on members. CS is built to meet it more cheaply and more credibly than a discretionary central bank, for reasons internal to its channel architecture. The claims in this section are computed on the framework's actual issuance engine (Paper 1, §3.3); the model validates CS's capacity to hold the anchor, while Section 7 validates the choice of a near-zero anchor.

Hitting zero does not disarm the demand instrument. In the unbundled channel structure, distribution and stabilization are separate instruments. The inflation-gap channel KI is the only over-the-line channel; setting its target to the flat price path holds inflation at zero. The growth dividend to citizens — K3 — is price-neutral and fires on real growth independently of the inflation setting. The split parameter κ_d governs how growth is shared between the indexed stake (K2) and non-indexed cash (K3), again independently of the price target.

Zero is safe for CS where it is dangerous for orthodoxy. The orthodox fear of low inflation is the zero lower bound: with demand carried by the policy interest rate, a near-zero price path risks a liquidity trap. CS carries demand through dividends (K1–K3), not through the price of credit, so the lower bound is not a binding constraint. Price stability does not disarm CS's demand instrument.

Low variance is structural, not discretionary. The predictability the layer requires is not a matter of central-bank resolve. With money creation decoupled from bank lending and governed by rule, the dominant source of inflation variance in the incumbent system — pro-cyclical credit creation — is absent. CS's low variance is a property of its plumbing, which is what makes its anchor credible to the external layer.

6. Holding Zero Costs Nothing on the Dividend

The sharpest way to see the decoupling is in dollars, computed on the real engine at launch calibration. The citizen consumer dividend is $K3 = \kappa_d \cdot (g_r \cdot M2 - K1_{agg})$: it is drawn from the same real-growth budget that funds the locked stake K2, so it does not depend on the inflation setting at all. At the launch calibration the growth budget is $g_r \cdot M2 = 2\% \times \$22,366B = \$447B$, and with K1 = $9B aggregate the dividend pool is $438B. The table shows the dividend across three inflation regimes for $\kappa_d = 0.5$, against a bundled counterfactual in which the only way to pay a citizen dividend is issuance above the growth line.

Table 1. Citizen dividend (K3) versus bundled-system dividend, by inflation regime.

Inflation regime	CS citizen dividend K3 (κ_d = 0.5)	Bundled-system dividend
Mode A — mild deflation (≈ −1.86%)	$219B (K2 to floors $219B; total on the $447B growth line)	not applicable — bundled systems do not run deflation by design
Mode B — price stability (0%)	$219B	$0B — no above-line issuance to distribute
Mode C — mild inflation (+2%)	$219B	$447B (= 2% × M2)
(+3% for comparison)	$219B	$671B

The CS dividend is $219B at every inflation setting — flat — because it is funded by growth, not by inflation. The bundled dividend collapses to $0 at price stability: to pay citizens it must inflate. This is the result that makes the common-zero anchor cheap for CS: choosing zero costs the citizen dividend nothing, because distribution (K3) and the price level (KI) are separate instruments drawing on separate budgets. In a bundled architecture, choosing price stability would mean forgoing distribution; in CS it does not.

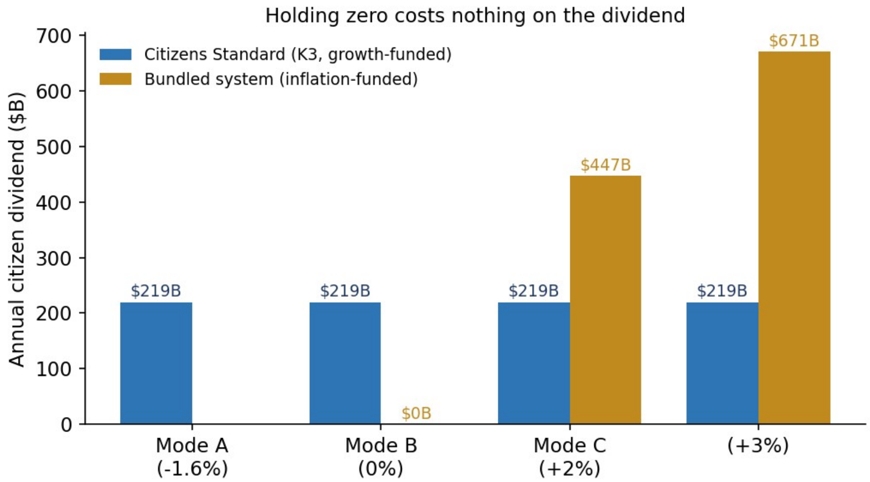

Figure 1. The citizen dividend (K3) is flat across inflation regimes in CS because it is growth-funded; a bundled system's dividend collapses to zero at price stability. Launch calibration; source: cs_engine.py.

7. Evidence

The external-layer results below are produced on the settlement layer's exact formula using a realistic wage-indexation lag (wages catch up; the real-rate gap is constant, not compounding), confirmed by an adversarial red-team and a stress-test across alternative wage models and realized variance. The model runs heterogeneous economies (productivity growth 3.0% and 0.5% — within the observed cross-country range) over a 40-year horizon. The parameter ranges span the empirically documented spread of cross-country productivity growth and of wage stickiness (staggered/Calvo wage-setting in the sense of Taylor, 1980 and Calvo, 1983).

Table 2. Summary of interoperability results.

#	Result	Finding
1	Layer is real-neutral	Nominal drift offsets the inflation differential exactly; deflated real rate = 1.000000 (machine precision).
2	Full indexation	Inflation cancels in H; heterogeneous-level and common-level real-rate paths are identical. The level is irrelevant when wages fully index.
3	Heterogeneous levels distort modestly and constantly	Under realistic indexation a 4-point level gap distorts the bilateral real rate by ≈ −3.9% (1-yr wage lag), −11.2% (3-yr), −18.0% (5-yr) — constant over time, not compounding. The −37.6% figure arises only if wages never catch up (a pessimistic bound).
4	Only common zero is robust to heterogeneous stickiness	With different stickiness across members (1- vs 4-year lag), a common +3% leaves ≈ −8.5%, a common +1% ≈ −2.9%, a common 0% exactly 0. Zero is robust to every stickiness profile.
5	Cost of a positive anchor	Excess price level on a high-productivity economy over 40 years: corridor +3% adds +226% /

		+388% / +632% versus a 0% / −1% / −2% baseline; a zero anchor adds +0% / +49% / +124%. The corridor is the costly choice at every baseline.
6	Differential variance, not level, drives volatility	Nominal-rate volatility ≈ 0.28%/yr at a large level gap with low variance; ≈ 2.79%/yr at high idiosyncratic variance; ≈ 0.56%/yr when shocks are correlated across members (ρ = 0.8). Level is irrelevant; correlation cancels in the differential.
7	Zero-robustness is general	At a common zero the inflation-driven distortion is ~0 under three wage processes (lag, partial-adjustment, Calvo) — clean for lag and Calvo; partial-adjustment leaves a ~1% residual that is a growth-tracking effect, present at any anchor and independent of inflation. Under a common +3% the residual is substantial in every process (lag −8.5%, partial-adjustment −7.6%, Calvo +5.4% / −5.2%); its sign tracks which member is the stickier one, so its direction cannot be known without each member's wage microstructure.

Results 1–2 establish that the level is irrelevant under ideal conditions. Result 3 sizes the real-world distortion honestly — modest and bounded, not the headline −38%. Result 4 is the robustness core: a common positive level does not eliminate the distortion under heterogeneous stickiness; only a common zero does. Result 5 shows zero is also the cheapest level at every baseline. Result 6 establishes that the always-binding requirement is the variance of the differential — predictability — which a level corridor does not address. Result 7 confirms the robustness is general across wage-setting models: at a common zero the inflation-driven distortion vanishes in all of them, while a common positive level leaves a residual whose magnitude is process-dependent and whose sign depends on which member is stickier — so under a corridor the bias cannot be signed in advance.

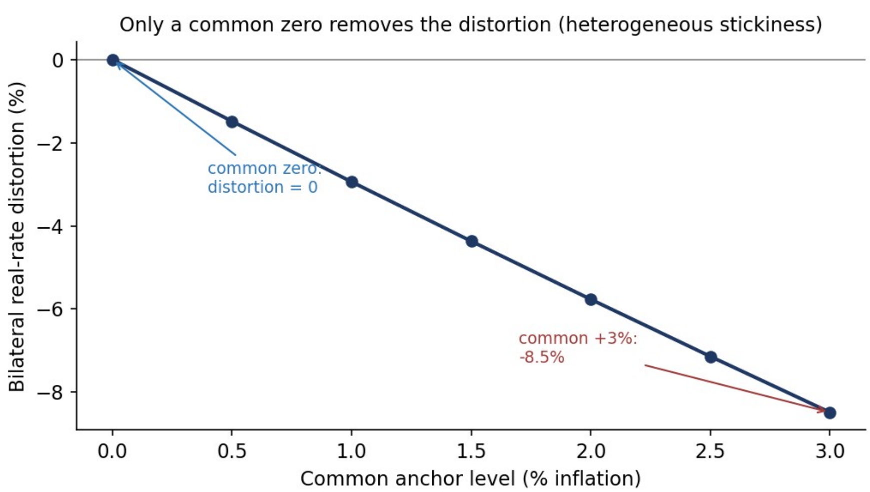

Figure 2. Bilateral real-rate distortion versus the common anchor level under heterogeneous wage stickiness: zero at a common zero, growing with the level. Source: equa_model_v3.py.

Interoperation reproduced inside the real engine. Coupling the CS issuance engine to the settlement layer reproduces the external-layer result inside the integrated domestic model: two members at the common-zero anchor settle on a clean bilateral rate reflecting only their productivity gap, while a member running a +3% corridor under sticky wages distorts the rate by −5.7%. The corridor distorts; zero does not — the same finding the standalone layer model produces.

8. The Load-Bearing Assumption: Independence Is Asymmetric

The decoupling of distribution from the price level (Section 6) is not unconditional. A red-team of the assumption finds it holds in one direction and must be defended in the other. When holding the anchor requires CS to add money — the design case, where the economy's natural drift is deflationary — independence is free: KI adds, K3 is untouched, and the dividend is unaffected. This two-sided behaviour, KI adding under deflationary drift while an inflationary drift is met by attenuating issuance rather than cutting the dividend, is formalized as a standing configuration in Mode Ω (Architecture, Section 8), with the mild-inflation attenuation rung specified as Tool 14c. When holding the anchor requires net contraction beyond KI's independent withdrawal capacity — an inflationary drift or shock — the decoupling can break, because defending the anchor would otherwise reach into the dividend. The resolution is to name the contractionary instrument explicitly (Section 9) so that KI never has to. With that instrument in place, independence holds in both directions; without it, only upward. Two further stress directions — a shared issuance ceiling and a small real cost of inflation — do not break the decoupling; both in fact reinforce the case for zero, since a positive corridor competes with the dividend under a cap and the dividend is maximized at zero.

9. The Contractionary Instrument: Retiring Money, Not Sterilizing It

The instrument that defends the anchor against an inflationary gap is the graduated surcharge specified in Paper 1, §10 (Tool 14b, the Anchor-Keeping Surcharge), which this paper's external-interoperability commitment activates. The design choice is to retire withdrawn money rather than absorb it into interest-bearing debt. The distinction is load-bearing, and the model makes it precise.

A bond-sterilization defense accumulates an interest-bearing stock governed by the carry condition. Over a 40-year persistent +3% drift the stock reaches roughly 56% of GDP when $r < g$, 80% when $r = g$, and 119% — and rising without bound — when $r > g$; the defense is sustainable only under positive carry. A graduated surcharge that retires the withdrawn money accumulates no stock at all: its cost is a bounded, non-compounding annual drag (here ≈ 2% of GDP at the modeled gap), and there is no $r < g$ sustainability condition because nothing is rolled over. A transient inflationary shock is therefore absorbed with the dividend left intact and the anchor defended. The honest limit is the same one stated for Tool 14b: a persistent real imbalance — production short of absorption — would require the surcharge to run continuously, and its bound is then the real and political tolerance for that standing drag, not a divergent stock. At that boundary the anchor yields to real adjustment. No monetary instrument neutralizes a real imbalance forever; what the retire-don't-sterilize choice buys is that the defense itself never becomes the instability.

10. Onboarding

A high-inflation legacy economy can join the anchor without a disinflationary collapse in distribution. Modeled on the real engine, a member entering at ~19% inflation converges to the common anchor over roughly a decade (≈11.6% at year 4, ≈1.7% at year 8, ≈0 by year 12) while the real citizen dividend remains continuous throughout — it grows with real output the entire way down. Disinflation under CS is therefore not paid for by suspending distribution, because the two are separate instruments; the onboarding path tightens the price level (via KI and, where needed, the surcharge) without touching the growth-funded dividend.

11. Propositions

Proposition 1 (Real-neutrality of the layer). Under a computed real-purchasing-power exchange layer with a relative-PPP inflation update, bilateral inflation differentials pass through to the nominal rate and leave the real rate unchanged. Inflation differentials cannot destabilize the real exchange rate through the layer.

Proposition 2 (Zero is the uniquely robust anchor). Under incomplete wage indexation, heterogeneous national inflation levels distort the bilateral real rate. A common level reduces the distortion, but a common positive level leaves a residual whenever members' wage stickiness differs; only a common zero anchor eliminates the inflation-driven distortion for every wage-adjustment process and every cross-country stickiness profile, because at zero inflation there is no inflation–wage gap to act on. Robustness, not merely cost, selects zero.

Proposition 3 (The anchor). The level that is simultaneously uniquely robust (Proposition 2), cost-minimizing at every baseline, and the non-arbitrary focal point is a common near-zero point, held with low variance of the differential. The interoperability anchor is a common, credible, low-variance, near-zero point — not a positive corridor.

Proposition 4 (Domestic feasibility). In the unbundled channel architecture the citizen dividend is funded from the real-growth budget and is invariant to the inflation setting; holding the common-zero anchor therefore costs the dividend nothing, and the anchor is defended by a money-retiring surcharge that accumulates no stock. A CS member can hold the anchor at no cost to its growth-distribution channel and without exposure to the zero lower bound.

Corollary. The binding multilateral requirement on a CS member is a predictable, low-variance price path at a common near-zero point — which CS meets structurally, demonstrated on its own issuance engine at launch calibration.

12. Limitations and Scope

Several scope limits are stated plainly. The external settlement layer is a proposed mechanism with no operating instance, so the external-layer results (Section 7) are a calibrated mechanism property, not an empirical forecast; the bilateral illustration uses two representative economies parameterized within empirically observed ranges rather than the microdata of named country pairs. The domestic engine reproduces the framework's published launch calibration, and its multi-decade comparisons are deterministic baselines; stochastic characterization is carried in the companion empirical paper. Under a partial-adjustment wage process the layer shows a small (~1%) residual even at a common zero, which is a productivity-wage-tracking effect present at any anchor and independent of inflation — reported for completeness, it does not bear on the zero-versus-positive comparison. The KI price-path determinacy and stability results are inherited from the Macroeconomic Model paper under transactional-circuit indexing. The Technical Appendix develops each of these in full.

13. Materials and Reproducibility

All quantitative claims in this paper are reproducible. A replication package accompanies it containing every model script, a one-command driver (run_all.py), a pinned environment (requirements.txt), and a results manifest mapping each numbered claim to its script and captured output. The domestic engine (cs_engine.py) reproduces the framework's published launch figures to the dollar; the external-layer scripts reproduce the Section 7 results; the structural scripts reproduce Sections 6 and 8–10. This paper's Technical Appendix (below) provides the formal model specification, proofs of Propositions 1–4, the full parameter table with sources, the external-layer calibration, and the robustness battery. Replication repository: github.com/Neo-Solon/Citizens-Standard.

Technical Appendix

External Interoperability and the Common Anchor — notation, model specification, proofs, calibration, robustness, and reproducibility.

Notation

Symbols used in this paper, with the section in which each is introduced.

Symbol	Meaning	Intro
$E(A,B)$	Bilateral settlement rate between monetary areas A and B: the ratio of labor-hours to buy the essentials basket in each.	§3
H_i	Labor-hours at the median wage required to buy the fixed essentials basket in economy i ($H = P/W$).	§3
P_i , W_i	Price level and median nominal wage index in economy i.	§3
π_i	Inflation rate in economy i; the differential $\pi_A - \pi_B$ is what the nominal rate passes through.	§3
Q_t	Real (productivity-adjusted) settlement rate: the nominal rate deflated by relative price levels.	App. B
g_i	Real labor-productivity growth in economy i.	§7
D_t	Distortion: the settlement rate relative to the productivity-implied benchmark, minus one.	§7
S_i	Member-specific wage-stickiness operator acting on the inflation path.	App. B
M2	Broad money supply; launch calibration $22,366B.	§5
g_r , n	Real output growth (2%) and population growth, governing the growth budget and KI.	§5
γ	Gap-closure rate of the KI inflation-gap channel.	§5
K1, K2, K3	Issuance channels: citizenship stake; growth routed to locked floors; price-neutral consumer dividend.	§5
KI	Inflation-gap channel — the only over-the-line issuer; defends the anchor by adding money.	§5
κ_d	Split parameter allocating the growth budget between locked floors (K2) and the dividend (K3).	§6
r , g	Interest rate and growth rate in the bond-sterilization carry condition ($r < g$ sustainability).	§9

A. Model Specification

A.1 The external settlement layer

CS specifies an external settlement layer of the computed real-purchasing-power class. The class was introduced by Serra's EQUA mechanism, on which CS builds; what follows is the layer's formal specification as used in CS.

For two monetary areas A and B, the layer publishes a rate equal to the ratio of labor-hours, at median wage, required to purchase a fixed constitutionally defined essentials basket:

$$E(A,B) = H_A / H_B, \quad H_i = P_i / W_i$$

where P_i is the area's price level and W_i its median nominal wage index. The rate is updated by a relative purchasing-power-parity factor and a bounded real-trade-flow term that cannot override it:

$$E_t = E_0 \cdot \Pi_{s \le t} (1 + \pi_{A,s}) / (1 + \pi_{B,s}).$$

The real (productivity-adjusted) rate is the nominal rate deflated by relative price levels, $Q_t = E_t \cdot (P_{B,t}/P_{A,t})$. The benchmark productivity rate is bench = $\Pi(1+g_A)^{-1} / \Pi(1+g_B)^{-1}$, the rate implied by the productivity gap alone. The distortion reported throughout is $D_t = (H_{A,t}/H_{B,t})/\text{bench} - 1$.

A.2 The CS issuance engine

Implemented from Paper 1, §3.3 and Appendix B. New money flows through rule-bound channels, distributed equally per citizen:

- **K1 (citizenship):** 0.025 × (GDP per capita) per verified citizenship event.
- **K2 (growth stake, locked):** the residual of the growth-matched line after K1; in Mode B, $K2 = g_r \cdot M2 - K1_{agg}$.
- **K3 (consumer dividend):** $K3 = \kappa_d \cdot (g_r \cdot M2 - K1_{agg})$, drawn from the same growth budget as K2 (the split parameter κ_d allocates the growth budget between locked stake and spendable dividend).
- **KI (inflation-gap, price-level-path targeting):** $KI = \max(0, (\pi^* + g_r + \lambda \cdot gap) \cdot M2 - K1 - K2)$, where π^* is the inflation target, g_r total real growth, λ the gap-closure speed, and gap = ln(target_price_level(t−1)/actual_price_level(t−1)) the log price-path gap (population growth does not enter the aggregate rate); see Paper 1, §6.1.

The engine reproduces the published launch figures to the dollar (K1 $2,250/citizen, $9.0B aggregate; growth line $447B; K2 $438B; KI $443B in Mode C) — see cs_engine.py.

B. Propositions and Proofs

Proposition 1 — Real-neutrality of the layer

Statement. Under the relative-PPP update, bilateral inflation differentials pass through to the nominal rate and leave the real rate unchanged.

Proof. With price levels $P_{i,t} = P_{i,0} \cdot \Pi_{s \le t}(1+\pi_{i,s})$, substitute the nominal-rate law into $Q_t = E_t \cdot (P_{B,t}/P_{A,t})$:

$$Q\_t = E\_0 \cdot [\Pi(1+\pi\_A)/\Pi(1+\pi\_B)] \cdot [P\_\{B,0\}\Pi(1+\pi\_B)] / [P\_\{A,0\}\Pi(1+\pi\_A)] = E\_0 \cdot (P\_\{B,0\}/P\_\{A,0\}).$$

The inflation products cancel identically for every t, so Q_t = Q_0: the real rate is invariant to the inflation differential. ∎ (Numerically, the deflated real rate equals 1.000000 to machine precision — equa_redteam.py.)

Proposition 2 — Zero is the uniquely robust anchor

Statement. Under incomplete wage indexation, a common zero anchor eliminates the inflation-driven distortion for every wage-adjustment process and every cross-country stickiness profile; any common positive level leaves a residual whenever members' stickiness differs.

Proof. Write H_i = P_i/W_i and decompose the wage index into a productivity-tracking component (captured by, and cancelling against, the benchmark bench) and a price-catch-up component governed by a member-specific stickiness operator S_i acting on the inflation path. The distortion reduces to D = f(π; S_A, S_B), the differential price-catch-up gap. Two facts follow. (i) For a common zero path, π ≡ 0, so the operator has no price movement to lag: the catch-up gap is identically zero for every S_i, hence f(0; S_A,S_B) = 0 for all stickiness profiles. (ii) For a common positive level π > 0, the catch-up gap is nonzero whenever S_A ≠ S_B; its magnitude increases in π and its sign is determined by the stickiness ordering of A versus B, not by the level alone. Therefore zero is the unique common level for which the inflation-driven distortion vanishes across all stickiness profiles simultaneously. ∎

Confirmed numerically across three wage processes (lag, partial-adjustment, Calvo): at a common zero the inflation-driven distortion is ~0 in all; at a common +3% the residual is −8.5%, −7.6%, and +5.4%, and reverses to −5.2% in the Calvo case when the stickiness assignment is flipped — i.e. its sign tracks which member is stickier, as the proof states (equa_stress.py, S3). The partial-adjustment process shows a ~1% residual even at zero; this is a productivity-wage-tracking effect, present at any anchor and independent of inflation, and so does not bear on the zero-vs-positive comparison.

Proposition 3 — The anchor

Statement. The interoperability anchor is a common, credible, low-variance, near-zero point — the level that is simultaneously uniquely robust, cost-minimizing at every baseline, and the non-arbitrary focal point — not a positive corridor.

Proof. The proposition is the conjunction of three properties established above. (i) *Uniquely robust* — by Proposition 2, a common zero is the unique common level that eliminates the inflation-driven real-rate distortion across all wage-stickiness profiles simultaneously; any positive level leaves a stickiness-dependent residual. (ii) *Cost-minimizing* — by Proposition 4 and §6, the CS citizen dividend is growth-funded and invariant to the inflation target, so holding zero imposes no dividend cost, and the money-retiring surcharge that defends the anchor accumulates no stock; no common level meets the predictability standard at lower cost. (iii) *Non-arbitrary focal point* — zero is the unique candidate that fixes no discretionary parameter: any positive corridor (e.g. 2–4%) requires heterogeneous members to negotiate and maintain an arbitrary number, whereas zero is the Schelling point that coordinates them without a bargaining surface. Properties (i)–(iii) jointly single out the common, credible, low-variance, near-zero anchor. ∎

Proposition 4 — Domestic feasibility

Statement. Holding the common-zero anchor costs the CS citizen dividend nothing; in a bundled architecture it costs the entire dividend.

Proof. By the issuance specification, $K3 = \kappa_d \cdot (g_r \cdot M2 - K1_{agg})$ depends only on real growth and the citizenship aggregate; the inflation target enters issuance solely through the separate KI channel. Hence $\partial K3/\partial(\text{inflation target}) = 0$. In a bundled architecture the only money available to distribute as a dividend is issuance above the growth-matched line, which equals zero at price stability. Therefore the dividend is invariant to the price target in CS and collapses to zero at price stability in a bundled system. ∎ (Numerically: K3 = $219B at Modes A/B/C; bundled dividend $0 / $447B / $671B at 0% / +2% / +3% — cs_engine.py, fig. A.1.)

C. Parameters and Sources

Every parameter used in the paper, its value, and its source. Launch macro-aggregates are 2025-vintage public series as anchored in the CS Macroeconomic Model paper (Neo-Solon, 2026e, App. A.1).

Parameter	Value	Source
M2 money supply	$22,366B	FRED series M2SL (2025 vintage); Neo-Solon 2026e A.1
Nominal GDP	$30.8T	BEA national accounts (2025)
Population	341.8M	U.S. Census (2025)
GDP per capita	$90,000	Derived (GDP / population)
Real growth g_r	2.0%	Long-run trend assumption (CBO range); robustness in companion empirical paper
Population growth n	0.5%	Census projection (Table 5 base case)
K1 rate	2.5% of GDP/capita	Paper 1 §5 (constitutional calibration)
κ_d (growth split)	0–1 (0.5 illustrative)	Mode parameter, Paper 1 §3.3/§7.4
Surcharge ceiling	3% of M2 ($671B)	Paper 1 §10 (Tool 14)
φ_{liq} (pledgeable floor)	≈0.15 (locked)	Macro model §; banking separation analysis
κ_W (MPC, asset wealth)	0.03 (0.02–0.05)	Poterba 2000; Chodorow-Reich et al. 2021; Di Maggio et al. 2020
Productivity gap (A,B)	3.0% / 0.5%	Within observed cross-country range — Feenstra-Inklaar-Timmer 2015 (PWT); World Bank 2021
Wage stickiness range	1–5 yr lag; Calvo θ 0.4–0.7	Card-Hyslop 1997; Barattieri-Basu-Gottschalk 2014; Daly-Hobijn 2014 (≈12-mo mean duration)

D. Calibration of the External-Layer Parameters

The external-layer results are a property of the settlement layer's exact formula evaluated over empirically grounded parameter ranges; the layer has no real-world instance, so no backtest is possible and none is claimed.

Productivity dispersion. Long-run labor-productivity growth varies widely across economies — from near-stagnation in some mature economies to sustained catch-up growth above 3% in others (Feenstra, Inklaar & Timmer, 2015; World Bank Global Productivity, Dieppe ed., 2021). The 3.0%/0.5% pair used in the bilateral illustration represents a high-growth catch-up economy alongside a mature low-growth one, both inside the documented range.

Wage stickiness. Micro evidence places the average duration of nominal wages at roughly twelve months, with a hazard spike at one year and pronounced downward rigidity (Card & Hyslop, 1997; Barattieri, Basu & Gottschalk, 2014; Daly & Hobijn, 2014). The paper's indexation-lag range (1–5 years) and Calvo reset durations (≈1.7–3.3 periods) bracket this central value and its tail. The theoretical forms are the staggered-contract and random-reset models of Taylor (1980) and Calvo (1983).

E. Robustness and Sensitivity

The central claim — zero is the uniquely robust anchor — was subjected to an adversarial red-team (equa_redteam.py) and a stress battery (equa_stress.py). None overturned it; several sharpened it.

Table 3. Robustness checks on the indexation-distortion result.

Check	Result
Is the −38% distortion real?	No — it assumes wages never catch up. Realistic indexation lags give a small, constant offset (−3.9% to −18%). The honest figure is bounded and non-compounding.
Does any common level cancel?	No — only a common zero. Under heterogeneous stickiness a common +3% leaves a residual (−8.5%); a common zero leaves 0.
Is the cost ratio baseline-specific?	The '8×' is specific to the −1% baseline; the robust statement is that a +3% corridor adds 200–400% cumulative excess price level while zero adds ~0.
Does variance beat level if shocks correlate?	Yes — it is the variance of the differential that matters; correlation (ρ=0.8) cuts it from 2.79 to 0.56 %/yr. Level remains irrelevant.
Survive realized variance, not just zero mean?	Yes — at a zero mean the systematic distortion stays ~0 as realized variance rises to 2%/yr; variance adds noise, not bias.
Cost of a deflation buffer?	Bounded — lifting the mean to +0.5% costs ≈1.5% residual, scaling linearly with the mean. The band is a clean knob.
Specific to one wage model?	No — holds under lag, partial-adjustment, and Calvo. At zero all give ~0; at +3% the residual's sign tracks which member is stickier.

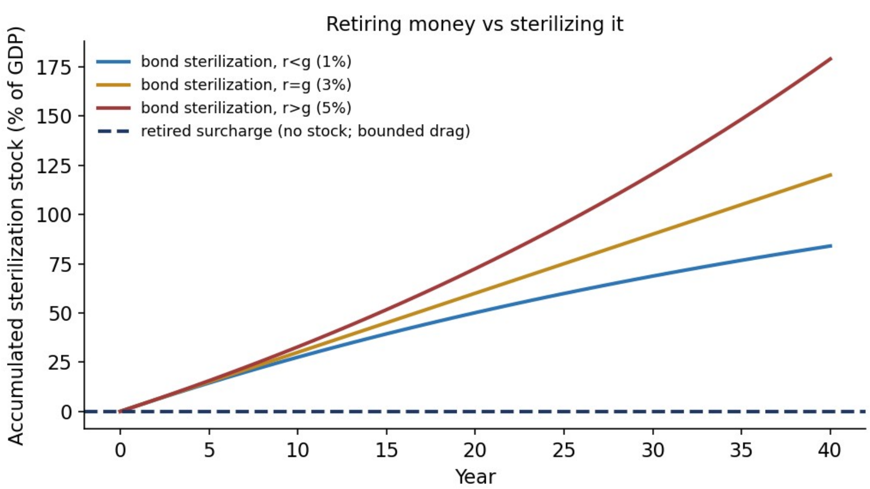

Figure A.1. Defending the anchor: bond sterilization accumulates an interest-bearing stock governed by r vs g, while a retiring surcharge accumulates none. Source: cs_contraction_compare.py.

F. Limitations and Scope

- The external settlement layer is a proposed mechanism with no operating instance; Section 7's results are a calibrated mechanism property, not an empirical forecast. The bilateral illustration uses two representative economies parameterized within empirical ranges, not the price and wage microdata of named country pairs.
- The domestic engine reproduces the framework's published launch calibration; the 40-year projections are deterministic baselines. Stochastic characterization (return, inflation, and growth resampling) is carried in the companion empirical paper, not here.
- Under a partial-adjustment wage process the layer shows a ~1% residual even at a common zero. This is a productivity-wage-tracking effect, present at any anchor and independent of inflation; it does not affect the zero-versus-positive comparison but is reported for completeness.
- The KI price-path determinacy and stability results are inherited from the Macroeconomic Model paper under transactional-circuit (M^T) indexing; this paper does not re-derive them.

G. Reproducibility

All results reproduce via the replication package: pip install -r requirements.txt, then python run_all.py. Every script is deterministic (fixed seeds). RESULTS_manifest.md maps each numbered claim to its script and captured output line. Environment: Python 3.12, numpy 2.4, matplotlib 3.10. The domestic engine (cs_engine.py) reproduces the published launch figures to the dollar; the external-layer scripts (equa_model_v3, equa_redteam, equa_stress) reproduce the Section 7 results table; the structural scripts (cs_channel_test, cs_independence_redteam, cs_contraction_compare, cs_sterilization_test) reproduce Sections 6, 8–10.

References

Barattieri, A., Basu, S., & Gottschalk, P. (2014). Some evidence on the importance of sticky wages. American Economic Journal: Macroeconomics, 6(1), 70–101.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. Journal of Monetary Economics, 12(3), 383–398.

Card, D., & Hyslop, D. (1997). Does inflation grease the wheels of the labor market? In Reducing Inflation (NBER), 71–122.

Chodorow-Reich, G., Nenov, P., & Simsek, A. (2021). Stock market wealth and the real economy. American Economic Review, 111(5), 1613–1657.

Daly, M. C., & Hobijn, B. (2014). Downward nominal wage rigidities bend the Phillips curve. Journal of Money, Credit and Banking, 46(S2), 51–93.

Di Maggio, M., Kermani, A., & Majlesi, K. (2020). Stock market returns and consumption. Journal of Finance, 75(6), 3175–3219.

Feenstra, R. C., Inklaar, R., & Timmer, M. P. (2015). The next generation of the Penn World Table. American Economic Review, 105(10), 3150–3182.

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.* SSRN Working Paper 6873798.

Neo-Solon (2026e). *The Citizens Standard: A Macroeconomic Model of a Two-Circuit Monetary System.* SSRN Working Paper 6939418.

Neo-Solon (2026f). *The Citizens Standard: Full-Reserve Banking and the Two-Circuit System.* SSRN Working Paper 6939498.

Replication code and data for all results: github.com/Neo-Solon/Citizens-Standard.

Poterba, J. M. (2000). Stock market wealth and consumption. Journal of Economic Perspectives, 14(2), 99–118.

Serra, D. (2026). Public Cash Money — Technical Framework, Ch. 6–7. publiccashmoney.com.

Taylor, J. B. (1980). Aggregate dynamics and staggered contracts. Journal of Political Economy, 88(1), 1–23.

U.S. Census Bureau (2025). National Population Estimates and Projections. census.gov.

World Bank (2021). Global Productivity: Trends, Drivers, and Policies (A. Dieppe, ed.).