
THE CITIZENS STANDARD

The Issuance Engine

Channel Mechanics, the Transactional-Money Locus, and Bounded Citizen Ownership

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: 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: Empirical Validation of the Transactional Aggregate (Neo-Solon, 2026j) · SSRN 6973298

The Citizens Standard: Governance and the Political Economy of the Parameters (Neo-Solon, 2026k) · SSRN 6973318

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 gives a complete and transparent account of the computational core of the Citizens Standard — the issuance engine — and then confronts the four places a skeptical economist is most likely to attack. The engine is shown to reduce to a single object: one growth-tied money budget per year, split across four channels (K1, K2, K3, KI), routed between two monetary circuits (a transactional circuit M^T that sets the price level, and an asset circuit M^a that builds a bounded ownership stake). Inflation and citizen ownership are not parameters but outputs, read off the back of the channel settings. We derive the price-stability locus $g \cdot M^T \approx \$229.7B$, show that all creation channels can be set to zero (yielding the fixed-supply, mildly deflationary limit), and formalise the structural buyer, whose ownership share ψ^* is bounded above by c/g and cannot run to the whole market. We then treat, honestly and with sensitivity analysis, the four contested points: the M^T decomposition, the 20% asset-to-goods spillover, the market-cap ownership assumptions, and the governance of the constitutional parameters. In each case the framework's structure is shown to be robust to plausible calibration error: the contested figures are empirical inputs with stated identification strategies and bounded effects, not load-bearing constants.

JEL classification: E31, E42, E51, E52, E58, G12

Keywords: Citizens Standard; issuance engine; money creation; two-circuit monetary system; price-stability locus; structural buyer; bounded ownership; sensitivity analysis

1. Introduction

The Citizens Standard is, at bottom, a rule for creating money and a rule for distributing it. Every macroeconomic claim the series makes — price stability under a standing dividend, a built savings floor, the compression of the return on capital toward the growth rate — is ultimately a claim about what one engine computes when its dials are set a certain way. This paper documents that engine in full, at the level of detail an economist needs to reproduce it, and then does the thing a serious proposal must do: it goes to the places the design is most exposed and meets them directly.

Two commitments organise what follows. First, *transparency*: every quantity below is exactly what the engine evaluates, channel by channel, with no step hidden in prose. Second, *falsifiability*: where a number is an empirical estimate rather than an identity — and three of them are — we say so, give the identification strategy, and report how far the conclusions move when the estimate is wrong. A framework that cannot survive its own sensitivity analysis does not deserve adoption; one that can should be judged on its structure, not on a contested decimal.

Sections 2–6 lay out the mechanics: the two-circuit architecture (§2), the four issuance channels and their formulae (§3), inflation as a derived quantity and the price-stability locus (§4), the zero-issuance limit (§5), and the index-versus-dividend distinction with the bounded structural buyer (§6). Section 7 — the heart of the paper — confronts the four anticipated lines of critique. Section 8 synthesises the robustness picture; §9 concludes. Throughout, the calibration is the United States: broad money $M2 \approx \$22.4T$, real growth $g \approx 2\%$, GDP $\approx \$30.8T$, population $\approx 342M$, and a domestic equity-index capitalisation of $\approx \$69T$.

2. Architecture: one budget, two circuits

The engine issues, at most, one growth-matched budget of new money each year, and the single analytical move that makes the rest tractable is to recognise that newly issued money does not all act on prices. It lands in one of two circuits. The *transactional circuit* (M^T) is money that circulates against currently-produced goods and services; by the quantity identity it, and only it, sets the goods price level:

$$P = M^T \cdot V / Y$$

The *asset circuit* (M^a) is money that is locked into the citizen savings floor and held as equity. It bids on existing assets, not on groceries, so it is inert with respect to goods prices except through a small, explicitly-modelled leakage (the spillover of §4). The decomposition is the whole trick: it is what allows a large gross issuance to be price-neutral, because most of it can be routed into M^a , where it builds wealth without chasing output. Holding this distinction, the modes stop looking like sleight of hand and become what they are — different routings of one budget between two circuits.

The transactional share is written $\mu = M^T/M2$, with a calibrated value $\mu \approx 0.5135$, so $M^T \approx \$11.5T$. This single number is the most consequential empirical input in the framework, and §7.1 is devoted to defending it.

3. The four issuance channels

Each year the engine evaluates four channels in a fixed order. The first sizes a per-citizen endowment; the second sizes the whole growth budget; the third routes that budget between the two circuits; the fourth, used only when inflation is deliberately wanted, issues above the growth line.

3.1 K1 — Citizenship

A one-time endowment paid into the floor for each new citizen, sized as a share k_1 of GDP per capita:

$$K_1 = k_1 \cdot (GDP / N) \cdot B$$

N is population and B the gross number of new citizens in the year. K_1 is paid first, off the top of the budget. The default is $k_1 = 2.5\%$; setting $k_1 = 0$ creates no endowment.

3.2 K2 — the growth budget

The throttle that sets how much new money exists. The full growth-matched budget is $M2 \cdot g$; the dial k_2 issues a chosen share of it:

$$G = k_2 \cdot M2 \cdot g$$

At $k_2 = 100\%$ the budget is the full \$447B; at 17.5% it is \approx \$78B; at 0% there is no budget at all. After K_1 is removed, the residual that the next dial routes is

$$R = \max(0, G - K_1)$$

3.3 K3 — the dividend share (κ_d): a router, not a spigot

This is the most misread channel, because it creates no money of its own. It *routes* the residual budget between the two circuits. A share κ_d is paid out as a spendable dividend (into M^T); the complement is locked into the floor (into M^a):

$$D = \kappa_d \cdot R \quad F = (1 - \kappa_d) \cdot R$$

The two shares always sum to one, so moving the dividend up moves the floor down by the identical dollars: *the total issued is unchanged*. K_3 is therefore price-relevant — it decides how much of the budget reaches M^T — but it is never a way to create additional money. $\kappa_d = 0$ is all-floor; $\kappa_d = 100\%$ is all-dividend.

3.4 KI — the inflation gap

The only channel that issues *above* the growth line, and so the only one that deliberately creates inflation. It is additive — piled on top of the growth budget, not carved from it — and lands in M^T :

$$KI = k_I \cdot M2$$

Setting $k_I = 0$ leaves no above-line issuance. The framework's inflationary mode (Mode C, §4) uses this channel to fund a visible monthly dividend at the cost of a target rate of inflation.

Table 1 collects the four channels.

Channel	Formula	What it does
K1	$k_1 \cdot (GDP/N) \cdot B$	Per-new-citizen endowment into the floor; paid first.
K2	$k_2 \cdot M2 \cdot g$	Sets the size of the whole growth budget.
K3 (κ_d)	$D = \kappa_d \cdot R$; $F = (1 - \kappa_d) \cdot R$	Routes the budget between spendable dividend and locked floor; total unchanged.
KI	$k_I \cdot M2$	Above-line issuance; the sole inflation channel.

Table 1. The four issuance channels. $R = \max(0, G - K_1)$ is the residual budget after the citizenship endowment.

4. Inflation as a derived quantity

The engine never asks for an inflation rate; it computes one. What matters is the flow that actually reaches the transactional circuit: the dividend D , the inflation-gap issuance KI , and the spillover L of asset-circuit money into goods demand. Writing A^* for the floor inflow (§6) and λ for the spillover fraction,

$$L = \lambda \cdot A^*, \quad \lambda = 0.20$$

the realised goods inflation is the transactional injection expressed against the transactional stock, net of real growth:

$$\pi = (D + KI + L) / M^T - g$$

If transactional money grows at exactly the real rate, the two terms cancel and prices are flat. This defines the *price-stability locus*: the system is price-neutral precisely when the injection equals $g \cdot M^T$.

$$D + KI + L = g \cdot M^T \approx \$229.7B$$

Locked floor money (the $K2/M^a$ share) barely appears, because it is not chasing goods — it enters π only through the modest spillover L . This is why a mode can issue a large gross budget and still hold prices flat: it routes the bulk into the asset circuit. Table 2 shows the four base modes as points in this space; each inflation figure can be checked by hand from the formula above.

Mode	k_1	k_2	κ_d	κ_l	Reaches M^T	Inflation
A deflation	2.5%	17.5%	0%	0	spillover only ($\approx \$16B$)	-1.86%
B stable	2.5%	100%	40%	0	div \$179B + spill $\approx \$54B$	0%
C inflation	2.5%	17.5%	0%	1.98%	KI \$443B ($\approx \$108/cit\cdot mo$)	+2.0%
D pure dividend	0	51.35%	100%	0	div \$230B (on the locus)	0%

Table 2. The four base modes. Modes B and D reach price stability by opposite routings — B locks 60% and pays 40%, D locks nothing and pays all — yet both place $\approx \$230B$ on the transactional locus.

5. The zero-issuance limit

Because every creation channel has an independent floor at zero, the engine contains the no-issuance economy as a genuine limiting case. Setting $k_1 = k_2 = \kappa_l = 0$ makes $K1 = 0$, $G = 0$ (hence R, D, F all zero), and $KI = 0$. No new money is created — not less, none. The system is not a printing press with the dial fixed on 'more'; zero is one of its admissible settings.

The consequence is the opposite of the printing-money intuition. With no new money but real output still growing at g , the same money stretches over more goods, and the inflation identity returns

$$\pi = 0 / M^T - g = -g \approx -2\%$$

A dollar gains roughly 2% a year. This is the fixed-supply, hard-money corner — the gold-standard outcome — recovered exactly as a limit of the same machine. Mode A (mild deflation) sits one step in from it; Modes B and D (stability) a few steps further; Mode C (mild inflation) one step beyond stable. The framework is therefore neither inflationary nor deflationary by nature: it spans the range, and the dials choose the point, including the point at which nothing is created at all.

6. Index versus dividend: the bounded structural buyer

The K3 routing decides more than ‘locked or spendable’; it decides whether a citizen’s share of new money becomes *ownership* or *income*. The floor share is not held as cash: it is used to buy equity in the domestic market index and held, making each citizen a part-owner of productive capital. The dividend share is spendable cash — consumed, gone once spent, building no stake. The structural buyer is the system standing in the market each year, buying the index on citizens’ behalf, with an annual flow

$$A^* = K_1 + (1 - \kappa_d) \cdot R$$

Measured against total index capitalisation M_{index} , this is a per-year bite $c = A^* / M_{index}$. At the default \$69T market, Mode B’s $A^* \approx \$272B$ is $c \approx 0.39\%$ of the entire market per year. Holding for dur years under growth g does not simply stack to $c \cdot dur$, because the market itself grows and floors are eventually drawn down; the realised collective ownership share is a growth-discounted accumulation,

$$\psi^* = c \cdot a(g, dur), \quad a(g, dur) = (1 - (1+g)^{-dur}) / g$$

with a hard ceiling that is the crux of feasibility: regardless of holding period,

$$\psi^* \leq c / g$$

At the defaults (40-year hold, 2% growth), Mode B settles near $\psi^* \approx 10\%$ of the index owned collectively, against a ceiling $c/g \approx 20\%$, leaving $\approx 90\%$ as tradable private float. Ownership *asymptotes*; it does not compound without limit. Mode D, by contrast, buys no index at all ($A^* = 0$). The economic significance of the ceiling is taken up in §7.3.

7. Where the framework is most exposed

Four points will draw the sharpest scrutiny. None is a structural flaw; each is a place where a number must be defended rather than asserted. The honest response is the same in every case: name what kind of object the number is, give the identification strategy, and report the sensitivity. Where the number is an empirical estimate, the test is not whether it is exactly right but whether the conclusions survive its being somewhat wrong.

7.1 The M^T decomposition (the largest vulnerability)

The entire inflation mechanism runs through $M^T \approx 51.35\%$ of $M2$. The fair objection is twofold: why that figure, and why should it be stable over time? We concede the centrality and answer in four parts.

First, what M^T is. It is the velocity-relevant, transaction-active share of broad money — balances actually turning over against current output — as opposed to the savings and near-money components of $M2$ that are held as stores of value. The quantity identity prices goods off the former, not the latter; the decomposition is standard monetary economics, not an invention of this framework.

Second, identification. The 51.35% figure should not rest on a single source, and does not need to. It can be triangulated from three independent routes: (i) the active-balance share implied by the $M1$ -versus- $M2$ composition; (ii) a velocity decomposition that attributes the measured fall in $M2$ velocity to a rising idle share; and (iii) the value back-solved from the observed price-stability locus itself. Convergence of these routes — not the precision of any one — is what should carry the

number, and a published confidence band ($M^T/M2 \in$ roughly 0.46–0.57) is more honest than a bare point estimate.

Third, and decisively, the rule is specified on *measured* M^T , not on the constant 0.5135. The framework issues at $g \cdot M^T$ for the M^T of the day, re-estimated on the same cadence as any other monetary aggregate. Payment technology and interest rates move the transactional share, but slowly; the rule re-anchors to it rather than freezing it. The 51.35% is today's best estimate, not a structural constant.

Fourth, the sensitivity is modest and bounded. Suppose the system issues the calibrated \$229.7B while the true transactional share differs from 51.35%. Table 3 shows the resulting miss: a six-percentage-point error in M^T — a large error for an aggregate — moves realised inflation by only about a quarter of a point, and in a direction the next re-estimation corrects.

True $M^T / M2$	M^T	Locus $g \cdot M^T$	Inflation if \$230B issued
45.0%	\$10.06T	\$201B	+0.28%
48.4%	\$10.81T	\$216B	+0.12%
51.35% (calibrated)	\$11.49T	\$230B	0.00%
54.4%	\$12.16T	\$243B	-0.11%
57.0%	\$12.75T	\$255B	-0.20%

Table 3. M^T sensitivity. Issuing the calibrated \$229.7B against a mis-estimated transactional share; the inflation miss over a ± 6 -point range stays within $\pm 0.3\%$.

7.2 The 20% spillover

The claim $L = 0.20 \cdot A^*$ invites the obvious question: why 0.20 and not 0.10 or 0.40? Again, four parts.

What it represents: the fraction of new asset-circuit (floor) money that leaks into goods demand — a wealth effect, the marginal propensity to consume out of newly-accrued floor wealth. It is the one channel through which locked savings touch goods prices, and it is why M^a is *nearly*, not perfectly, price-inert.

Its anchor: the empirical wealth-effect literature places the marginal propensity to consume out of wealth at single-digit cents per dollar for liquid holdings and lower for illiquid ones. The floor is locked and illiquid by construction, which pushes the leakage of each year's inflow toward the low end; 0.20 of the annual flow is a deliberately conservative — cautious-high — mid-estimate rather than an optimistic one, chosen so the price result is stress-tested against more spillover than is likely, not less.

Its sensitivity: Table 4 varies the fraction across 10–40% for Mode B. The whole range moves realised inflation by under three-quarters of a point, and — crucially — the dividend share κ_d is a free instrument that re-centres the result: a higher measured spillover is offset by routing marginally less to the dividend. The spillover sets a calibration, not an outcome.

Spillover λ	Dividend + spillover	Injection	Mode B inflation
10%	\$179B + \$27B	\$206B	-0.21%
20% (assumed)	\$179B + \$54B	\$233B	+0.03%

Spillover λ	Dividend + spillover	Injection	Mode B inflation
30%	\$179B + \$82B	\$261B	+0.27%
40%	\$179B + \$109B	\$288B	+0.51%

Table 4. Spillover sensitivity (Mode B). Across a plausible 10–40% range the inflation effect is bounded and offsettable by κ_d .

7.3 Market-cap ownership assumptions

The structural-buyer account rests on four moving parts an economist will press immediately: long-run market returns, retirement drawdown, market-cap growth, and citizen participation. The decisive point is that the design already contains the object that disciplines all four — the ceiling c/g and the realised share ψ^* — which is what prevents runaway ownership and is a material improvement over earlier, unbounded versions of the idea.

Drawdown and holding period: ψ^* is bracketed by two decumulation models (a fixed-duration annuity and a constant-hazard process) and is monotone in the holding period. Table 5 varies that period from 20 to 60 years; ψ^* moves only between 6% and 12%, and in every case sits below the c/g ceiling of $\approx 20\%$, leaving the large majority of the market as tradable float.

Holding period	ψ^* (realised)	Bracket	Tradable float
20 years	6.0%	5.6–6.4%	94.0%
30 years	8.1%	7.4–8.8%	91.9%
40 years (default)	9.8%	8.7–10.8%	90.2%
50 years	11.1%	9.8–12.4%	88.9%
60 years	12.2%	10.7–13.7%	87.8%

Table 5. Ownership sensitivity (Mode B). Across all plausible holding periods ψ^* stays bounded below the $c/g \approx 20\%$ ceiling; float never falls below $\approx 88\%$.

Market-cap growth is self-correcting: if the index grows faster than g , the per-year bite $c = A^*/M_{index}$ falls, and with it ψ^* — the buyer’s footprint shrinks exactly when the market it is buying gets larger. Participation enters linearly: aggregate ψ^* scales with take-up, so partial participation lowers collective ownership and raises float; the distribution of floors across participants is a separate, distributional question treated in the macro paper. Finally, the realised return on the floor is not a free-market rate but the general-equilibrium, attenuated marginal product at universal scale (the $r \rightarrow g$ compression); it is model-dependent, calibrated on a Cobb–Douglas / Solow–OLG structure, and the *qualitative* results — bounded ownership, preserved float, return compression — are robust across the parameter range tested, even where the absolute floor values are not pinned to the last dollar.

7.4 Governance of the constitutional parameters

A political economist will ask the right question: who sets k_1 , k_2 , κ_d , and k_I ? The engine treats them as *constitutional parameters* — chosen through a slow, public, statutory process and changeable only through a tiered, transparent procedure, never by quarter-to-quarter discretion. The design intent is precisely to remove the discretionary lever that is the usual object of monetary capture, and to replace it with rules that are hard to move and visible when moved.

We state plainly that the governance *mechanism* — the constitutional encoding, the amendment thresholds, the emergency-adjustment protocol, the anti-capture safeguards — is outside the

scope of this document, which is about what the engine computes once the parameters are set. It is developed in the statutory paper of the series, and the updatability and governance design in the architecture paper. A reader is right to ask the question here; the honest answer is that it is answered there, and that no exposition of the mechanics should be read as a claim about the politics of setting them. The two must stand or fall together, but they are different papers.

8. Robustness: what the sensitivities add up to

Read together, the three empirical inputs behave the same way. None is an identity; each has a stated identification strategy and a bounded effect. A six-point error in the transactional share moves inflation by about a quarter-point. The spillover, across a 10–40% range, moves Mode B by under three-quarters of a point and is offsettable by the dividend dial. The ownership share stays between 6% and 12% across every plausible holding period and can never exceed its c/g ceiling. In each case the rule is specified on the measured quantity and retuned to it, so a calibration error shifts the dial setting, not the existence of a price-stable, bounded-ownership configuration.

This is the ordinary terrain of applied macroeconomics, not a special fragility. The framework does not claim its three numbers are certain; it claims that its *structure* — two circuits, a growth-tied budget, a routed dividend, and a bounded structural buyer — survives reasonable uncertainty in them. That is the standard a serious proposal should be held to, and the one this engine is built to meet.

9. Conclusion

The engine is far simpler than its output suggests. One growth-tied budget is sized by two dials, routed by a third between a spendable dividend and a locked, index-buying floor, and optionally supplemented by a fourth that is the only source of inflation. Two circuits do the analytical work: money reaching the transactional circuit moves prices; money locked in the asset circuit builds a bounded ownership stake. Inflation and ownership are read off the back of the settings, not dialed in; the no-issuance economy is a genuine limit; and the structural buyer is bounded above by construction.

The places critics will press — the M^T decomposition, the spillover, the ownership assumptions, the governance of the parameters — are real, and we have tried to meet them in the spirit a sceptic deserves: not by denying that the numbers can be wrong, but by showing how far wrong they can be before the conclusions change, and by pointing to where each is identified and defended in full. The mechanics are in this paper; the deeper rigor — the M^T identification, the welfare-optimal dividend share, the realisable-return derivation, and the constitutional governance — is in the architecture, macro-model, structural-buyer, and statutory papers of the series. The engine is the operational core; this paper is the map to it.