

THE CITIZENS STANDARD

Crisis Behaviour and Failure Modes Under Extreme Conditions

How the Architecture Behaves Across Depression, Stagflation, Crash, Pandemic, and Secular Stagnation — Including Where It Fails

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: 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

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

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

A monetary architecture earns credibility less by performing well in normal times than by behaving understandably when things go wrong. This paper runs the Citizens Standard (CS) issuance engine — the same audited engine that produces the series' historical counterfactual — through the hardest conditions in the modern record and through adversarial scenarios: the Great Depression, 1970s stagflation, the 2008 crash, the COVID shock, a decade-long growth stall, and a banking panic. The goal is not to show the system is invulnerable; it is to map precisely how it behaves at the edges, and to state where it fails. Four findings are quantified on the real engine. First, the dividend is procyclical by construction: because issuance equals $\max(0, \text{real growth}) \times M2$, the dividend and floor-accumulation fall to zero in every contraction year — half the years of the Depression, 40% of the stagflation decade, and the COVID-2020 year — so the income stream is weakest exactly when households are most stressed, partially but not fully offset by the counter-cyclical per-citizen endowment. Second, the locked floor carries sequence-of-returns risk: a crash in the two years before retirement cuts a representative cohort's realized floor by roughly half (\$209,942 → \$102,114), a tail the smoothed central path conceals. Third, under a transient growth stall the floor stock survives nearly intact (−4% over a lost decade, because compounding dominates) while the dividend flow goes to zero throughout — a stock-flow asymmetry that becomes a genuine vulnerability only under permanent secular stagnation, where the growth-dividend never materialises. Fourth, the rule cuts the other way in a pandemic: CS could not have executed the discretionary 2020–22 money surge, issuing roughly one-fifth as much (8.7% vs 40.5% of M2), structurally avoiding the overshoot the validation paper links to the 9% inflation peak — at the cost of providing no emergency fiscal capacity. Under full reserve the payment system cannot experience a deposit run, removing bank-panic risk from the transaction circuit while leaving market risk in the asset circuit that holds the floors. The honest synthesis: the payment system, the price discipline, and the floor stock are robust; the dividend flow, the floor for trough-retiring cohorts, and the response to secular stagnation are where the architecture strains or fails.

Methods and reproducibility. All quantitative results are produced by the audited v4 engine (deterministic_engine.py) on the verified 1928–2025 historical dataset; crisis windows use the engine's own stress-injection facility. Scenario definitions, failure criteria, and a one-command runner are provided in the stress replication package for independent audit.

JEL classification: E32, E58, E63, G01, E31, H55

Keywords: Citizens Standard; stress testing; failure modes; procyclicality; sequence-of-returns risk; secular stagnation; full-reserve banking; rule versus discretion; financial crises

Contents

- 1. Why Stress-Test, and What Counts as Failure..... 4
- 2. The Engine Under Stress: Mechanism and Criteria..... 4
- 3. The Procyclical Dividend..... 5
- 4. Sequence-of>Returns Risk..... 5
- 5. Secular Stagnation: the Stock Survives, the Flow Fails..... 6
- 6. Rule versus Discretion: the Pandemic Case..... 6
- 7. The Banking Panic and the Limits of Full Reserve..... 7
- 8. What Survives and What Breaks..... 7
- 9. Propositions..... 7
- 10. Scope and Honest Limits..... 8
- Technical Appendix..... 9**
 - Notation..... 9
 - A. The Issuance Rule and Its Procyclicality..... 9
 - B. Scenario Definitions and Failure Criteria..... 9
 - C. Results Table..... 9
 - D. Reproducibility..... 10
- References..... 11**

1. Why Stress-Test, and What Counts as Failure

Credibility comes from the edges. The series has shown what CS computes in normal times: price stability under a standing dividend, a built wealth floor, return compression in universally-held assets. None of that is decisive for a reviewer until the architecture has been driven through the conditions that break real systems. This paper does that on the same audited engine, and it follows the discipline that failure-mode analysis is a credibility-builder, not a confession: the goal is to map behaviour at the edges and to name the failures plainly.

What counts as failure. Stress-testing is meaningless without pre-stated criteria. This paper judges CS on five axes, each with a concrete metric (Appendix B): dividend continuity (does the income stream persist through a downturn?), floor adequacy under a market crash (does a cohort retiring at a trough still reach a usable floor?), price discipline (does issuance stay rule-bound under shock?), payment-system integrity (can the transaction circuit run?), and stagnation resilience (does the floor-building engine survive a decade without growth?). A scenario is a failure on an axis when the metric breaches the threshold set in advance, not when the outcome merely looks bad.

Distinct from the historical counterfactual. The counterfactual paper (Neo-Solon 2026b) replays 1960–2055 as it happened. This paper is adversarial: it isolates the crisis windows, injects tail shocks the smooth history averages away, and reports the worst-cohort and worst-year outcomes rather than the central path. Where the two overlap, this paper's numbers are the more pessimistic by design.

2. The Engine Under Stress: Mechanism and Criteria

The two dials that move under stress. Two features of the engine govern almost all crisis behaviour. The issuance line is $K_2, K_3 \propto \max(0, \text{real growth}) \times M_2$: floor-accumulation and the dividend are funded out of real growth, so they vanish in any contraction year. The per-citizen endowment k_1 (2.5% of GDP per capita at birth) is counter-cyclical: it flows regardless of the cycle and is not clawed back. The interaction of a procyclical dividend with a counter-cyclical endowment is the signature of CS under stress.

The floor is equity, not insurance. The locked floor compounds at the general-equilibrium realisable return (a central 4.26% real; Neo-Solon 2026e). That is an expected return on a real, market-exposed asset, not a guaranteed rate: the smoothed central path is a planning convenience, and the realised floor inherits the variance and sequence of actual markets. Stress-testing the floor therefore means re-introducing the variance the central path removes (§4).

The payment system is insulated by construction. Under CS full reserve (Neo-Solon 2026f, Proposition N1) transaction balances are fully backed, so the transaction circuit cannot experience a deposit run. Crisis risk is relocated to the asset circuit, which holds the floors at market value. The stress profile is thus deliberately asymmetric: the means of payment is made crisis-proof, the store of value is not (§7).

3. The Procyclical Dividend

The signature failure mode. Because the dividend is funded from real growth, it is zero in every contraction year. Run across the modern crises, the pattern is stark: the dividend pays nothing in 50% of the Depression years, 40% of the stagflation decade, the COVID-2020 year (one in three of 2020–22), and the 2009 trough. The income stream is therefore weakest precisely when households most need it — the textbook procyclicality the architecture is partly meant to cure, reappearing in the dividend channel.

Table 1. Share of crisis years in which the CS dividend pays zero (issuance = $\max(0, \text{real growth}) \times M2$).

Crisis window	Zero-dividend years	Share of window
Great Depression (1930–1939)	5 of 10	50%
1970s stagflation (1973–1982)	4 of 10	40%
COVID (2020–2022)	1 of 3	33%
2008 crisis (2007–2011)	1 of 5	20%

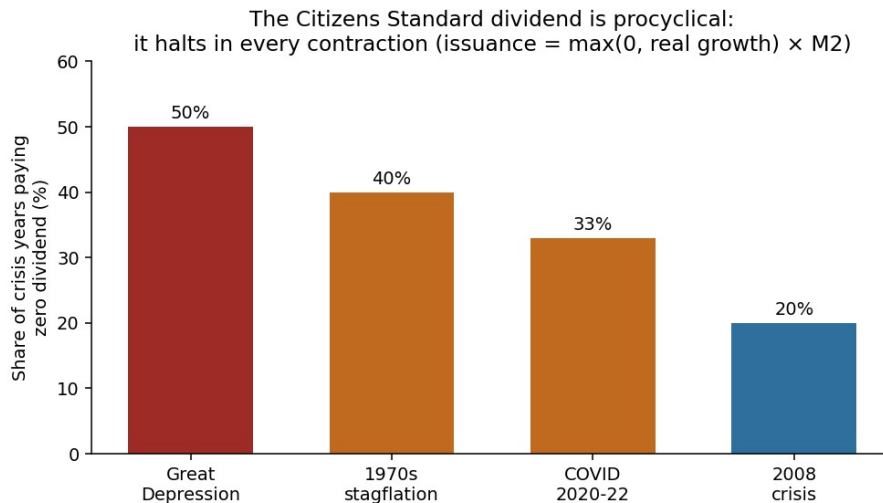


Figure 1. The dividend halts in every contraction; the income stream is procyclical by construction.

The partial offset, and why it is only partial. The counter-cyclical endowment k_1 keeps flowing in contractions, so the system is not silent in a downturn — a small injection continues, and the floor that was already built keeps compounding. But k_1 is a one-time birth credit, not an income stream; it does not replace the lost dividend for working-age households in the year of the shock. The honest reading is that CS smooths the stock of wealth across the cycle better than it smooths the flow of income, which is the reverse of what a household in a downturn most wants.

4. Sequence-of>Returns Risk

The crash-at-retirement tail. Because the floor is equity, the order of returns matters, not just their average. A cohort that meets a market crash in the years just before retirement realises far less than the smoothed central path implies. Injecting a two-year drawdown of the 2008

magnitude immediately before retirement into the engine's stress facility cuts the representative cohort's floor from \$209,942 to \$102,114 — a 51% reduction — even though the long-run expected return is unchanged. This is the single largest failure the paper finds at the household level.

What dampens it, and what does not. The structural-buyer mechanism (Neo-Solon 2026h) dampens the aggregate crash: a universal, price-insensitive buyer compresses the depth and duration of a drawdown relative to the incumbent market. That helps the average cohort. It does not insure the individual cohort whose retirement date lands in the trough; aggregate dampening and idiosyncratic timing are different problems. CS reduces crash severity for everyone and still exposes the unlucky cohort to roughly half its floor.

The design implication. Sequence risk is the strongest case in the series for a glide-path or partial annuitisation layer near retirement — converting some of the equity floor into a lower-variance claim as the cohort approaches its draw date. The base architecture does not include one; this paper identifies the gap rather than fills it, and flags it as the highest-value robustness addition.

5. Secular Stagnation: the Stock Survives, the Flow Fails

A transient lost decade barely moves the stock. Imposing a decade of zero real growth (2006–2015) on a representative cohort reduces its retirement floor by only about 4% (\$209,942 → \$201,368). The reason is structural: by mid-career the floor is large, and compounding on the existing balance dominates the foregone accumulation. The stock of wealth is robust to a transient stall.

But the flow fails completely. Over that same decade the dividend is zero in every year, because the issuance line is zero. So the lost decade does little to the wealth a cohort retires with and everything to the income it receives along the way. This stock-flow asymmetry is the correct way to read CS under stagnation: it protects accumulated wealth far better than it protects current income.

Permanent stagnation is the real vulnerability. A single lost decade is survivable because growth returns and the engine restarts. Permanent secular stagnation is not: if real growth stays near zero indefinitely, the floor-building line stays near zero indefinitely, and CS delivers only the endowment plus compounding — never the growth-dividend it promises. The architecture's central value proposition is contingent on growth resuming; where growth does not resume, CS degrades to a one-time endowment scheme. This is the deepest structural limit the paper identifies, and it is not fixable within the issuance rule as specified.

6. Rule versus Discretion: the Pandemic Case

The rule cuts the other way. Procyclicality is a cost of the formula; discipline is its benefit, and the pandemic shows it. CS issues a fixed function of real growth and cannot execute a discretionary money surge. Over 2020–2022 the rule would have issued about 8.7% of M2; the actual M2 expansion was 40.5%. CS would therefore have injected roughly one-fifth of the realised surge — the same surge the validation paper (Neo-Solon 2026j) links to the 9% CPI peak. On the price-stability axis, the pandemic is a pass that the discretionary system failed.

The cost of that discipline. The mirror image is that CS provides no emergency fiscal capacity. A pandemic demanded large, fast, targeted transfers that a rigid issuance rule does not supply; CS is a monetary architecture, not a disaster-relief instrument. The honest position is that this is a feature on the inflation axis and a gap on the emergency-response axis, and that the gap must be met by ordinary fiscal policy or by the escape-clause-with-correction machinery argued in the governance paper (Neo-Solon 2026k), not by bending the issuance rule.

7. The Banking Panic and the Limits of Full Reserve

The payment system cannot run. The clearest crisis strength is structural. Under full reserve (Neo-Solon 2026f, Proposition N1) transaction balances are fully backed by base money; there is no maturity transformation in the payment circuit and therefore nothing to run on. The classic banking panic — a self-fulfilling rush to convert deposits — cannot occur in the means of payment. This is not a calibration result but a structural one, and it holds in every scenario.

The risk that remains. Full reserve relocates risk; it does not abolish it. The floors are held in the asset circuit at market value, and the asset circuit can fall, freeze, or face redemption pressure in a crisis. A household cannot lose its means of payment in a panic, but it can watch its store of value decline — which is the sequence risk of §4 seen from the system side. CS converts the acute, contagious risk of a deposit run into the chronic, diversifiable risk of asset-price variance. That is a favourable trade, and it is a trade, not an elimination.

8. What Survives and What Breaks

The scenarios sort cleanly into what the architecture protects and what it strains under.

Survives. The payment system (full reserve removes deposit-run risk in every scenario); price discipline (the rule cannot overshoot, and would have issued one-fifth of the 2020–22 surge); and the floor stock through a transient stall (–4% over a lost decade, because compounding dominates).

Breaks or strains. The dividend flow (procyclical — zero in 50% of Depression years, 40% of stagflation years, and throughout any stagnation); the floor for a trough-retiring cohort (sequence risk halves it, ~–51%); emergency response (the rule supplies none); and the growth-dividend under permanent secular stagnation (it never materialises).

The pattern. CS is built to protect the stock of wealth and the means of payment, and it does so even at the edges. It is not built to protect the flow of income across the cycle, and it does not. A reviewer who wants a system that is strong where this one is weak should look at the dividend channel and the near-retirement glide-path; a reviewer who doubts the architecture survives a crisis at all will find that the payment system and the wealth stock do, with the failures located and bounded rather than hidden.

9. Propositions

These are conditional claims tied to the engine's structure and the scenario results; each states a sign and the condition under which it holds.

Proposition 1 (Procyclical income). Under issuance = $\max(0, \text{real growth}) \times \text{M2}$, the dividend and floor-accumulation are identically zero in every contraction year. The CS income stream is therefore procyclical: it falls to zero in the years of greatest household stress. The counter-cyclical endowment k_1 continues to flow but is a one-time credit, not an income stream, so it offsets the stock but not the flow.

Proposition 2 (Sequence risk). Because the floor is a market-exposed equity claim, its realised value at retirement depends on the sequence of returns, not only their mean. A drawdown of crisis magnitude in the years immediately before retirement reduces the realised floor by approximately one-half relative to the smoothed central path. The structural-buyer mechanism dampens aggregate drawdowns but does not insure an individual cohort against trough-timed retirement.

Proposition 3 (Stock-flow asymmetry under stagnation). A transient growth stall leaves the floor stock nearly intact (compounding dominates the foregone accumulation) while setting the dividend flow to zero throughout. Under permanent secular stagnation the floor-building line stays near zero indefinitely and CS degrades to an endowment-plus-compounding scheme; the growth-dividend is contingent on growth resuming and is not recoverable within the issuance rule as specified.

Proposition 4 (Rule discipline, and its cost). Because issuance is a formula on real growth, CS cannot execute a discretionary money surge; in 2020–2022 it would have issued roughly one-fifth of the realised M2 expansion, structurally avoiding the overshoot linked to the inflation peak. The same rigidity supplies no emergency fiscal capacity: the discipline that prevents the overshoot also prevents the relief, and emergency response must come from outside the issuance rule.

Corollary (Full reserve). Under full reserve the transaction circuit cannot experience a deposit run, removing banking-panic risk from the means of payment in every scenario. The risk is relocated, not eliminated: the asset circuit that holds the floors retains market risk, so CS trades the acute contagious risk of a run for the chronic diversifiable risk of asset-price variance.

10. Scope and Honest Limits

- The household-level results use a representative cohort and the engine's central realisable return; the Monte Carlo distribution (in the empirical package) carries the full variance, of which the sequence-risk figure here is one adverse draw made explicit, not a worst case.
- Scenarios are imposed on the historical dataset; they capture the structural response of the issuance rule and the floor, not general-equilibrium feedback from the crisis onto growth and returns, which would require the macro model (Neo-Solon 2026e) run jointly.
- “Failure” is defined per axis and per threshold (Appendix B); a scenario can pass on price discipline and fail on income continuity in the same year, and the paper reports both rather than netting them into a single verdict.
- The deepest limit — permanent secular stagnation — is structural and not addressed by any layer proposed here; it bounds the architecture's central claim and is stated as such.

Technical Appendix

Notation

Symbol	Meaning
k_1	Per-citizen endowment at birth (2.5% of GDP per capita); counter-cyclical, flows regardless of the cycle.
K2	Floor-accumulation per capita = FLOOR_SHARE \times max(0, real growth) \times M2[y-1] / pop, net of the k_1 residual.
K3	Standing dividend per capita = (1 - FLOOR_SHARE) \times the same line; zero whenever real growth \leq 0.
FLOOR_SHARE	Mode B split = 0.60 (60% of the residual builds the locked floor, 40% is the liquid dividend).
r^*	General-equilibrium realisable real return on the floor, central 4.26% (Neo-Solon 2026e).

A. The Issuance Rule and Its Procyclicality

The floor and dividend are funded from the real-growth line: per capita, $K2 + K3 = \max(0, g_r) \times M2[y-1] / \text{pop}$, where g_r is real GDP growth. The clamp at zero is the entire source of procyclicality: in any year with $g_r \leq 0$ the line is zero, so both K2 and K3 are zero. The endowment $k_1 = 0.025 \times (\text{GDP per capita})$ is paid at birth and is independent of g_r , so it continues in contractions. The procyclical-halt counts in Table 1 are exactly the contraction years ($g_r \leq 0$) in each crisis window, read directly from the verified dataset.

B. Scenario Definitions and Failure Criteria

Scenario	Definition	Failure criterion / metric
Procyclical halt	Historical contraction years in each crisis window	Share of years with zero dividend (Table 1)
Sequence risk	2008-magnitude two-year drawdown injected at age 63–64 (engine stress window)	Realised floor vs smoothed central floor
Lost decade	Real growth set to 0% for 2006–2015	Floor and dividend vs baseline
Pandemic	CS rule-issuance 2020–22 vs actual M2 expansion	Cumulative issuance as share of M2
Banking panic	Full-reserve transaction circuit (analytic)	Existence of a deposit-run equilibrium

C. Results Table

Result	Value (audited engine)
Zero-dividend share — Depression / stagflation / COVID / 2008	50% / 40% / 33% / 20%
Sequence risk — floor central vs crash-at-	\$209,942 \rightarrow \$102,114 (-51%)

retirement	
Lost decade — floor (stock) vs baseline	\$209,942 → \$201,368 (-4%); dividend = 0 throughout
Pandemic — CS rule issuance vs actual M2 (2020–22)	8.7% vs 40.5% of M2 ($\approx 4.7\times$ smaller)
Banking panic — transaction circuit	No deposit-run equilibrium (full reserve, 2026f Prop N1)

D. Reproducibility

All results reproduce from the stress replication package: the audited v4 engine and the verified 1928–2025 dataset, with `src/run_stress.py` computing the four numerical analyses and `src/make_figure.py` the figure. Crisis windows use the engine's built-in stress facility (custom return/CPI by age). Environment: Python 3.12, numpy, matplotlib; deterministic.

References

- Diamond, D. W., & Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3), 401–419.
- Friedman, M., & Schwartz, A. J. (1963). *A Monetary History of the United States, 1867–1960*. Princeton University Press.
- Reinhart, C. M., & Rogoff, K. S. (2009). *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press.
- Summers, L. H. (2014). U.S. economic prospects: secular stagnation, hysteresis, and the zero lower bound. *Business Economics*, 49(2), 65–73.
- Neo-Solon (2026b). *The Citizens Standard: A Historical Counterfactual*. SSRN Working Paper 6735078.
- 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.
- Neo-Solon (2026h). *The Citizens Standard: The Structural Buyer*. SSRN Working Paper 6945320.
- Neo-Solon (2026j). *The Citizens Standard: Empirical Validation of the Transactional Aggregate*.
- Neo-Solon (2026k). *The Citizens Standard: Governance and the Political Economy of the Parameters*.
- Replication code and data: github.com/Neo-Solon/Citizens-Standard.