
THE CITIZENS STANDARD

One Model, Many Systems

A Constitutional Monetary Architecture



Neo-Solon

Neo-Solon@hotmail.com

WORKING PAPER — JUNE 2026

COMPANION PAPERS

- 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: 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: 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

We hold these truths to be self-evident, that all men are created equal, that they are endowed by their Creator with certain unalienable Rights, that among these are Life, Liberty and the pursuit of Happiness. That to secure these rights, Governments are instituted among Men, deriving their just powers from the consent of the governed. That whenever any Form of Government becomes destructive of these ends, it is the Right of the People to alter or to abolish it, and to institute new Government, laying its foundation on such principles and organizing its powers in such form, as to them shall seem most likely to effect their Safety and Happiness.

Declaration of Independence, 1776

The Monetary Corollary

The right asserted in that passage is general. It turns on one question, whether a form of government serves the safety and happiness of the governed, and it exempts nothing, least of all the machinery of money. A monetary system is a governing power, arguably the most consequential a state delegates: its design determines who receives newly created purchasing power first, whose holdings appreciate and whose erode as the price level moves, and what share of future output is committed to servicing debt before that output exists. Powers of that magnitude are legitimate, in the founders' own logic, only so far as they serve the governed.

That the present architecture was established lawfully is not a defense against changing it; it is precisely what places it under the principle. Legislatures chartered the central banks, authorized the borrowing, and wrote the statutes under which money is issued, which is to say the system is a governing power in the fullest sense: created by government, and governing in turn, deciding who gains from newly created money and who pays for it. When such a power ceases to serve the safety and happiness of the governed, the people retain the right to alter or abolish it.

These consequences were not chosen at a ballot box. They follow from a design, and a design, unlike a law of nature, can be examined and altered.

This is not a separation of one people from another. It is the reconstitution of a single governing power, the authority to create money, upon the principle the founders named: that when a form of government becomes destructive of its ends, the people may institute a new one, organized in such form as shall seem most likely to effect their safety and happiness.

What follows is that architecture.

Abstract

Money in the United States is created at the discretion of unelected committees and of private banks, with no constitutional limit on its quantity and no rule governing its distribution. New money enters the economy hierarchically — reaching financial institutions first and ordinary households only later, after prices have already adjusted — a structure that has concentrated the gains from money creation and eroded the dollar's purchasing power by more than ninety-six percent since 1913. This paper proposes a constitutional alternative, the Citizens Standard, that addresses the problem at its two roots: who decides how much money is created, and who receives it.

The framework replaces committee discretion with a non-discretionary formula anchored to objective measures of real economic growth, and it distributes every newly created dollar equally to citizens rather than to institutions. Most of that issuance takes the form of a locked, individually owned stake in the productive economy — the Stable Floor — so that new money enters as citizens' ownership rather than as anyone's debt. A parallel reform separates the payment system from credit: transaction balances are fully reserved and constitutionally protected, so the payment system cannot fail when the credit cycle does, while banks continue to lend as private intermediaries.

Issuance flows through four rule-bound channels — a citizenship deposit at birth (K1), a growth-matched annual deposit (K2), an optional inflation-gap dividend (KI), and an optional growth-indexed consumer dividend (K3) — whose calibration lets a society constitutionally select its inflation regime: mild deflation with rising real wages, true price stability, or a modest citizen dividend. K3 shares K2's real-growth budget rather than adding to it, so a society can route a chosen share of each year's growth to citizens as current spendable income while holding the same price path. The paper specifies this architecture, distinguishes the unamendable core (the Model) from the configurable regimes it can host (the Systems), and situates the framework among existing monetary-reform and broad-capital-ownership proposals. This selection is meaningful only because the chosen regime is the unique determinate equilibrium: the money-quantity anchor of the two-circuit model pins the price level so that the targeted rate is actually delivered rather than being one of many self-fulfilling outcomes (Neo-Solon 2026e, Proposition 7).

Four base Systems are presented as illustrative configurations spanning two axes: the inflation-deflation spectrum, and the form in which new money reaches citizens — a locked Stable Floor versus an immediately spendable dividend. Mode A targets mild deflation of approximately 1.9 percent annually. K1 (citizenship) and K2 (growth) channels are active, routing into locked Stable Floor accounts via the capital markets channel. Mode A produces structurally rising real wages and is the appropriate choice for societies that prioritize purchasing power preservation. Mode B issues at the full real-growth-matched rate on a fixed 60/40 split. For an economy whose transaction-active money share sits near one-half it produces approximately zero annual drift; away from that structure the fixed split drifts (Section 5.x), and guaranteed price stability in any monetary structure is the role of Mode Ω (Section 8). K1 and K2 together fund the

circulating pool at the full real-growth-matched rate, producing a stable price level; the growth-matched budget is then split, sixty percent building a substantial locked Stable Floor and forty percent paid to citizens as a standing, immediately spendable dividend (K3) — a large capital stake and a current dividend together, both at zero inflation. Mode C targets approximately 2 percent inflation, activating the inflation channel — KI — that distributes new money equally to all citizens as an unlocked, immediately spendable citizen dividend of approximately \$108 per month at launch at a flat issuance rate of about 1.98 percent of M2 (3.85 percent of the transactional circuit M^T , the rate consistent with the 2 percent target); because that rate is a constant share of a growing transactional circuit, the dollar dividend scales gently with the economy, passing roughly \$115 per month within about two years and rising thereafter. Mode C is the configuration appropriate for societies that prefer current visible benefit during working life over long-term real-wealth compounding in the Stable Floor. Mode D shares Mode B's issuance regime — approximately stable near a one-half transaction ratio — but, unlike Mode B, builds no locked floor: it pays the entire real-growth-matched budget directly to citizens as a standing dividend (the dividend share κ_d is set to one hundred percent and the K1 citizenship channel is inactive), holding the price level flat because the dividend reallocates within the issuance budget rather than adding to it. Mode D leaves the capital decision wholly to the citizen — the dividend can be saved or privately invested at the roughly 6.7 percent unattenuated return that Mode D, deepening no aggregate capital stock, leaves intact — and suits societies that favor maximal current liquidity and individual portfolio choice over a universal locked stake. That this issuance registers as roughly two percent in consumer prices rather than being diluted across all broad money presupposes the two circuits remain separate — the separation whose determinacy consequence is established in Neo-Solon (2026e, Proposition 7) and whose institutional enforcement is the subject of the full-reserve banking paper (2026f).

These four configurations occupy two independent axes. The first is the price-level regime each targets: deflationary (Mode A), price-stable (Modes B and D), or mildly inflationary (Mode C). The second is the form in which citizen seigniorage is delivered: as a locked, indexed Stable Floor (Mode A, the indexed configuration); as a balanced split of floor and spendable dividend (Mode B, the balanced configuration); as a floor supplemented by an inflation-funded dividend (Mode C, the mixed configuration); or as a pure spendable dividend with no locked stake (Mode D, the distributed configuration). The two price-stable Modes, B and D, share a price level but sit at opposite ends of the second axis — B builds the largest locked floor of any base Mode, while D builds none, trading the universal capital stake for maximal current liquidity and individual portfolio choice. In short: A — deflationary/indexed, B — stable/balanced, C — inflationary/mixed, D — stable/distributed.

Beyond the four base Modes, the architecture supports any constitutional parameterization that preserves the Model's load-bearing properties. Mode Ω , introduced above and developed in Section 8, is the framework's flagship configuration: a two-sided stabilizing rule that solves for exact price stability in any economy, holding the price level flat in both directions rather than committing to a fixed split. It is the

cleanest demonstration of the central claim, that a single engine can guarantee a stable dollar as a matter of rule rather than discretion. The architecture's wider extensibility — that any coherent monetary regime within the Model's bounds can be constitutionally ratified — is developed through the further configurations of Sections 7.5 and 9. Section 10 presents Mode T, a transition-only configuration in the same illustrative spirit: it retires pre-existing sovereign debt through a dedicated channel, KT, while funding citizen Stable Floors at a full-rate, price-stable calibration, and then lands automatically in a permanent stable steady state once the debt stabilizes within its operational band.

All Systems use the same banking architecture, the same emergency toolkit, the same Market Exit, and the same supermajority amendment process. A society ratifying the Citizens Standard chooses, at ratification or by subsequent supermajority amendment, which Mode to operate under. The choice is reversible by the same supermajority process. The framework treats Mode selection as a legitimate constitutional choice rather than a technical optimization: a society that values rising real wages may prefer Mode A; a society that values nominal contract stability, with a current dividend at zero inflation, may prefer Mode B; a society that values a larger explicit dividend funded by mild inflation may prefer Mode C; a society that wants adaptive demographic and productivity governors may ratify a configuration like Mode Λ or design its own. The Model accommodates all of these because each preserves its load-bearing properties.

JEL classification: D63, E42, E51, E52, E58, E61

Keywords: Citizens Standard; constitutional monetary architecture; rule-based money creation; sovereign money; universal capital endowment; monetary regime design

Contents

1. The Problem with Discretionary Money Creation.....	11
1.1 How Money Is Currently Created.....	11
1.2 What Discretionary Creation Produces.....	11
1.3 What the Current System Gets Right.....	12
1.4 The Question This Paper Answers.....	12
2. The Three Governing Variables.....	12
2.1 The Variables.....	13
2.2 Why Existing Anchors Fail.....	13
2.3 Is New Money Actually Necessary?.....	14
3. The Mechanical Design — Model and System.....	15
3.1 The Model.....	15
3.2 The System.....	16
3.3 The Issuance Channels.....	16
3.4 Two Pools, Four Channels.....	17
3.5 Adaptive Smoothing and the Composite Productivity Index.....	18
3.6 Citizenship and K1.....	18
3.7 The Stable Floor Account.....	19
4. Mode A — The Deflationary System.....	21
4.1 Parameters.....	21
4.2 Outcomes.....	21
4.3 The Deflation-Debt Question.....	22
4.4 When Mode A Is Appropriate.....	22
4.5 The Hoarding Question and the Demurrage Backstop.....	22
5. Mode B — The Split-Dividend System.....	22
5.1 Parameters.....	23
5.2 Outcomes.....	23
5.3 When Mode B Is Appropriate.....	24
6. Mode C — The Inflationary System.....	25
6.1 The KI Formula.....	25
6.2 Parameters at Launch.....	26
6.3 How New Money Enters the System in Mode C.....	27
6.4 The KI Lock — and Why There Isn't One.....	27
6.5 Outcomes.....	27
6.6 Why This Is Not Universal Basic Income.....	28
6.7 The Cantillon Argument in Mode C.....	29
6.8 When Mode C Is Appropriate.....	29
7. Cross-Mode Comparison.....	29
7.1 At Launch.....	29
7.2 Tradeoff Summary.....	32
7.3 Choosing a Mode.....	32
7.4 Mode D — The Distributed System.....	33

7.5 Beyond the Four Base Modes.....	34
8. Mode Ω — The Price-Stability Configuration.....	33
8.1 The Two-Sided Stabilizing Rule.....	33
8.2 The Stabilizing Split Is Structural, Not Growth-Dependent.....	34
8.3 Domain of Validity and the Contraction Hand-Off.....	34
8.4 Relation to Mode B, and a Rejected Alternative.....	35
8.5 The Ownership Dial.....	35
9. Mode Λ — An Adaptive Configuration.....	36
9.1 Core Goal.....	37
9.2 K1 — Citizenship Endowment (Adaptive).....	37
9.3 K2 — Growth Dividend (Adaptive).....	38
9.4 K1 — Supplemental Citizen Stream (Conditional).....	38
9.5 Additional Built-in Features.....	38
9.6 Simulation Results — 2025-Born Citizen at Age 65 (in 2025 dollars).....	39
9.7 What Mode Λ Demonstrates.....	40
9.8 Constitutional Guardrails.....	41
10. Mode T — The Transition Configuration.....	40
10.1 Mode T-stable — The Automatic Steady State.....	41
11. Banking Architecture.....	44
11.1 The Payment-Credit Separation.....	45
11.2 How Credit Is Created.....	45
11.3 The Transition Lending Facility.....	45
11.4 Mode-Specific Banking Considerations.....	45
12. Emergency Tools and Crisis Response.....	46
12.1 Failure Mode 1 — Demand Collapse.....	46
12.2 Failure Mode 2 — Inflation Surge.....	47
12.3 Failure Mode 3 — Banking Liquidity.....	48
12.4 Failure Mode 4 — Credit Cycle.....	48
12.5 Cross-Mode Tools.....	48
12.6 Mode-Specific Tool Calibration.....	49
12.7 Total Crisis Response Capacity.....	49
12.8 Emergency Tool Mapping for Custom Modes.....	50
13. The Market Exit.....	50
13.1 What the Exit Provides.....	50
13.2 Exit Destinations.....	51
13.3 Exit Conditions.....	51
13.4 Strategic Effect.....	51
14. Honest Limitations.....	51
14.1 Limitations Common to All Modes.....	51
14.2 Mode A-Specific Limitations.....	52
14.3 Mode B-Specific Limitations.....	53
14.4 Mode C-Specific Limitations.....	53
14.5 Mode Λ -Specific Limitations.....	53
14.6 Limitations Not Solved by Mode Choice.....	54

14.7 The Equity Transmission Architecture as Structural Feature.....	54
15. Governance.....	56
15.1 The Constitutional Protocol.....	56
15.2 The Three-Tier Updatability Framework.....	56
15.3 Mode Selection as Constitutional Choice.....	57
15.4 Mode Stability and Why Modes Do Not Auto-Rotate.....	57
15.5 Statistical Independence.....	58
15.6 Identity Verification.....	59
16. System Comparisons.....	59
16.1 Comparison with Current System.....	59
16.2 Comparison with MMT.....	59
16.3 Comparison with Bitcoin.....	60
16.4 Comparison with UBI Proposals.....	60
16.5 The Citizens Standard’s Distinctive Position.....	60
17. International Applicability.....	60
17.1 Requirements for Any State.....	60
17.2 Mode Choice and National Context.....	61
17.3 Reserve Currency Implications.....	61
17.4 The Adaptable Core.....	61
18. Long-Run Considerations.....	61
18.1 Intra-Mode A Cycling.....	61
18.2 Inter-Mode Constitutional Change.....	62
18.3 Structural Risks Across All Modes.....	62
18.4 The Constitution Principle.....	62
18.5 A Complementary Fiscal Option: Land Value Taxation.....	63
18.6 Payout Neutrality and the Equity-Transmission Safeguard.....	63
19. Situating the Citizens Standard in the Literature.....	66
19.1 Rules versus Discretion, and the Constitutional-Money Tradition.....	66
19.2 Free Banking: The Decentralized Alternative.....	66
19.3 Full-Reserve Banking and Payment-Credit Separation.....	67
19.4 Broad Capital Ownership and the Social Dividend.....	67
19.5 Modern Monetary Theory as Contrast.....	68
19.6 The Gap the Framework Occupies.....	68
20. Conclusion.....	68
Glossary.....	69
References.....	73
Appendix A: Mode Interaction Simulations.....	77
A.1 Constitutional Mode Transitions.....	77
A.2 Generational Mode Rotation.....	78
A.3 Mid-Crisis Mode Transition.....	79
A.4 Figure Construction and Provenance.....	80
Appendix B: Interactive Engine.....	81
B.1 Availability.....	81
B.2 What the Engine Exposes.....	81

B.3 What the Engine Outputs.....	81
B.4 Honest Caveats.....	82
B.5 Open Release.....	82
Technical Appendix.....	83
A.1 Launch Parameters and Shared Variables.....	83
A.2 The Issuance Channels.....	84
A.3 Stable Floor Accumulation Model.....	87
A.4 Cross-Mode Outcomes — Full Parameter Table.....	89
A.5 The Composite Productivity Index.....	90
A.6 Emergency Toolkit — Ceiling Summary.....	91
A.7 Full Parameter Table.....	92
A.8 Mode Λ Governor Derivation.....	94

1. The Problem with Discretionary Money Creation

The United States dollar has lost more than 96 percent of its purchasing power since the Federal Reserve was established in 1913. What cost one dollar in 1913 costs approximately thirty-one dollars and fifty cents today. This is not an anomaly or a policy failure. It is the predictable, structural output of a monetary system in which money is created at the discretion of unelected committees with no constitutional constraint on the quantity, timing, or distribution of new money.

1.1 How Money Is Currently Created

In the fractional reserve system, commercial banks create money whenever they issue loans. The Bank of England confirmed in its 2014 Quarterly Bulletin that banks create new deposits — new money — at the moment of lending, not by transferring existing deposits. In the United States, the large majority of the money supply exists as commercial bank deposits created through this process, with state-issued physical currency making up only a small share — on the order of a tenth of M2. The Federal Reserve, separately, creates base money through open market operations, asset purchases, and direct credit facilities. Neither mechanism has a constitutional limit. Neither distributes new money equally. Both distribute it hierarchically — to financial institutions first, to the broader economy later, after prices have already adjusted.

This is the Cantillon Effect, named for the eighteenth-century economist Richard Cantillon: whoever receives new money first benefits at full purchasing power; whoever receives it last absorbs the inflation it caused. In the current system, the first recipients are always banks. The last recipients are always wage earners and savers.

This is more than an inflation problem. The current system produces a debt-based institutional order. Money enters the economy as someone's liability — either a bank loan that must be repaid with interest, or a Treasury bond that taxpayers must service. Every dollar in circulation traces back to a debt that someone owes someone else. In a debt-based society, the institutions that issue debt hold the lever that determines economic conditions. An equity-based civil society inverts this relationship: new money enters as citizens' equity — claims of ownership in the productive economy, held individually, not as anyone's liability.

1.2 What Discretionary Creation Produces

Between 2020 and 2022, the Federal Reserve expanded M2 by approximately 40 percent in two years. The result was 11.1 percent consumer price inflation in June 2022, the highest in forty years. To correct this, the Fed raised interest rates from near zero to over 5 percent in eighteen months, triggering a housing market freeze, the failure of Silicon Valley Bank and other regional institutions, and the destruction of real wages for three consecutive years. The cure was caused by the disease.

The pattern is not new. The 1970s stagflation lasted a decade because the Fed validated oil supply shocks with monetary expansion. The 2008 financial crisis was inflated by years of artificially cheap money. The national debt reached approximately \$39.2 trillion by mid-2026, with annual net interest payments approaching \$1 trillion, because government has been able to borrow against a currency it can expand at will.

1.3 What the Current System Gets Right

Intellectual honesty requires acknowledging what discretionary monetary authority genuinely accomplishes. The Federal Reserve's response to the 2008 crisis prevented a complete global payment system collapse. Its response to COVID-19, while inflationary, produced the fastest recession recovery in modern American history. The liquidity function — ensuring that payment systems do not seize during panics — is genuinely valuable. The Citizens Standard does not eliminate this function. It reduces the discretionary steering function while preserving the liquidity function through bounded, rules-based emergency tools described in Section 12.

1.4 The Question This Paper Answers

If discretionary monetary creation is the disease, what should replace it? The framework begins from a single premise: the value created when new money is issued belongs to the citizens of the polity equally — as owners, not as recipients of a discretionary transfer — rather than accruing, as it now does, to whichever institutions stand nearest the point of issuance. What remains open is not whether citizens should receive that value, but in what form and on what schedule. Different answers produce different societies. A frozen money supply produces structurally rising real wages. A growth-matched money supply produces price stability. A formula-targeted modest inflation produces a citizen dividend funded by citizen seigniorage. Each is defensible. Each implies tradeoffs. The Citizens Standard's argument is that the choice among them is properly a constitutional decision made by citizens through supermajority process — not a technical optimization made by committee. The framework provides the architecture that makes any such choice coherent and durable. The remainder of the paper specifies that architecture and presents illustrative Systems it can host; the full menu of valid Systems is bounded by the Model's load-bearing properties (Section 7.5) rather than by the number presented here.

2. The Three Governing Variables

Every monetary framework must answer one question: what determines how much money should exist? The current system answers: whatever a committee decides. The gold standard answered: however much metal exists in vaults. An absolutely fixed supply system answers: a predetermined, unchanging amount. The Citizens Standard answers with three variables working in coordination.

2.1 The Variables

Population is the trigger. Verified additions to the citizenry — births registered and naturalizations completed — initiate K1 issuance events. K1 is the citizenship channel: each qualifying event triggers a one-time deposit, calibrated as a fixed fraction of GDP per capita, into the new citizen's locked Stable Floor account. The Stable Floor is the constitutionally protected equity account every citizen holds from birth or naturalization until age 65 — its balance, held as total-market index shares, represents the citizen's guaranteed ownership stake in the productive economy. Population is directly countable, slowly and predictably changing, and internationally verifiable through public records.

GDP is the calibrator. K1 is calibrated as a fraction of GDP per capita at the moment of citizenship. K2 is the growth channel: calibrated annually as a fraction of measured real economic growth, it distributes new money in equal per-citizen amounts to all living citizens' Stable Floor accounts each month. Both auto-scale with the economy and remain meaningful across decades and centuries.

The price level is the inflation regulator. In Mode A, the price level feeds back into K1 through a deflation adjustment; in Mode B, K2 is matched to real growth so the price level holds approximately stable for economies near a one-half transaction ratio, without a deflation adjustment. In Mode C, the price level is the primary calibrator of K1 magnitude — the gap between actual and target inflation directly determines how much new money K1 creates each year.

2.2 Why Existing Anchors Fail

A monetary anchor must satisfy three conditions: it must be measurable, it must scale with the real economy, and it must be resistant to manipulation by the issuing authority. Gold satisfies the first and third but fails the second; pure GDP rules satisfy the first and second but fail the third. The Citizens Standard's anchor — citizenship plus a Composite Productivity Index plus a constitutionally specified price-level path — satisfies all three.

Gold is governed by geology, not economics. When the American population nearly tripled from 1870 to 1920, gold supply grew approximately 30 percent. The relationship between metal in vaults and economic output is essentially random; gold cannot anchor a monetary system to the productive activity it is meant to facilitate.

GDP-growth rules, most prominently associated with Milton Friedman's 1960 k-percent proposal, improve on gold by tying monetary expansion to economic output. But pure GDP anchors fail in two ways. First, methodological: GDP is a single statistic produced by a single agency on a quarterly revision cycle, creating both timing fragility and single-point-of-manipulation risk. Second, procyclical: GDP falls in recessions, so a strict GDP-anchored issuance rule contracts the money supply precisely when monetary contraction is most damaging.

The Citizens Standard's anchor is constructed to address both failures. The Composite Productivity Index combines five measures, each produced by a different federal agency, each updated on a different cycle: real GDP per worker (BEA), industrial

electricity consumption (EIA), freight ton-miles (BTS), total factor productivity (BLS), and port and rail throughput (Census and Association of American Railroads). The geometric mean of these five measures, each normalized to a common index base, produces a single productivity number that no single agency can manipulate, and that responds smoothly to actual economic activity rather than tracking quarterly GDP revisions. Four of the five inputs are independently auditable by foreign governments — a property no current reserve currency offers, and one that materially strengthens the dollar's case as the world's reserve currency under any Mode (Section 19.3).

The Composite Productivity Index is structurally load-bearing throughout the framework: K2 calibration depends on it in all Modes, KI calibration in Mode C is partly driven by it, and several emergency tools trigger on it. The framework's reliance on these statistical agencies makes statistical-agency independence the single most important governance commitment in the entire architecture. Section 15.5 develops this argument in detail.

Independence of the producing agencies is necessary but not sufficient. The formula is only as trustworthy as the integrity of the numbers fed into it, and a measured input can be silently revised after the fact even by an independent agency. The framework therefore imposes a measurement-integrity requirement on every load-bearing input — the Composite Productivity Index and the price-level series that drives KI and the emergency triggers. Each such input, together with the raw source data and the transformation code that produces it, must be published through a tamper-evident, independently reproducible pipeline: every release cryptographically committed and time-stamped, every revision permanently recorded and diffable rather than overwritten, and the entire chain verifiable by any third party — including foreign governments — without trusting the publisher. The requirement is on the property, not on any one technology. A public append-only ledger (a distributed ledger or blockchain) is one way to obtain it; a transparent, open-source, reproducible pipeline with cryptographic hashing, public timestamping, and independent replication is another. The specific implementation is therefore an operational matter (Tier 3), while the property itself is structural. What a tamper-evident layer cannot do is adjudicate methodology: it faithfully records whatever is committed to it, so the integrity of the inputs still rests on agency independence and on the publicly specified construction of each index. The two commitments are complements — independence governs how a number is produced, tamper-evidence governs that the produced number cannot be quietly changed. As with the privacy-preserving identity infrastructure, a measurement layer at this standard does not yet exist at national scale; building it is a three-to-five-year prerequisite, and the standard set here is the target architecture rather than a description of present capability.

2.3 Is New Money Actually Necessary?

This is the most fundamental question in the framework and deserves a direct answer: no, new money is not strictly necessary. New people earn their share of the existing money supply through work. Money circulates — one dollar can pay a million wages over time. A growing population sharing a slow-growing or fixed money supply produces

mild deflation as more people divide a pool that expands more slowly than the economy. This is what Mode A produces by design — not because circulating M2 is literally frozen, but because the only path by which new money reaches circulation in Mode A is the capital markets channel, which produces approximately 0.35 percent of M2 per year, well below real economic growth.

New money creation is a policy choice. The Modes of the Citizens Standard make that choice differently. Mode A creates new money only at the rate needed to fund Stable Floors, accepting deflation as a feature. Mode B creates new money at the full real-growth-matched rate on a fixed 60/40 split; this holds prices approximately stable near a one-half transaction ratio. Mode C activates the KI direct deposit channel at a magnitude calibrated to produce modest, predictable inflation, directing that citizen seigniorage to citizens as a visible monthly dividend. Importantly, KI's inflation effect is a function of its calibrated magnitude relative to real economic growth — not an inherent property of the channel itself. A society could run all three channels and still land in deflationary or stable territory if total issuance remains below real growth. Adaptive configurations like Mode A vary their creation rate in response to demographic and productivity conditions, activating KI conditionally. None of these is inherently correct. Each is a coherent constitutional choice, and the architecture is designed so that the choice itself is the citizens' to make and to revisit.

3. The Mechanical Design — Model and System

3.1 The Model

The Model is the unamendable architecture. It consists of seven load-bearing properties, each of which is preserved across all Modes.

- Dual-circuit separation. A circulating pool prices wages, commerce, and goods. A Stable Floor pool holds locked equity. The two pools intersect only through the capital markets channel, never through direct transfer.
- Citizen-anchored issuance. New money is created through verified civic events — citizenship, growth, or inflation gap — never through institutional transactions, government borrowing, or committee judgment.
- Equal per-citizen distribution at issuance. Whatever new money is created, it is distributed in equal shares to citizens, not hierarchically to institutions. This eliminates internal Cantillon advantage by construction.
- Equity-based ownership. The Stable Floor pool holds claims on productive businesses, not stored dollars. KI (in Mode C and conditionally in Mode A) routes through direct citizen deposits, but K1 and K2 always route through the equity index.
- Separated banking. Transaction accounts are full-reserve and constitutionally protected. Term deposits are at-risk investment claims that explicitly carry credit risk. The payment system cannot fail when the credit system does.

- Constitutional rules over committee discretion. All issuance is formula-based. Emergency tools are bounded, rules-triggered, and automatically reversing. The FDCA (Federal Digital Currency Authority) implements rules; it does not set policy.
- Supermajority amendability with citizen-only voting. Each verified citizen has exactly one vote. There is no wealth-weighted voting, no institutional proxy, and no emergency suspension of the amendment requirement.

These seven properties are Tier 1. They define the Model. They cannot be removed by amendment without abandoning the framework.

3.2 The System

The System is the specific parameterization a society chooses at ratification. The Citizens Standard presents four illustrative System configurations — four base Modes — each a complete, internally coherent application of the Model. The framework can host other valid Systems beyond these three; the choice among them is bounded only by the Model's load-bearing requirements (Section 7.5). Section 9 presents Mode Λ as one concrete example of an extended configuration.

Mode selection is a Tier 2 constitutional choice: it requires a 67 percent supermajority to ratify and a 67 percent supermajority to change, with a mandatory 90-day deliberation period. A society can ratify the Model and then debate which Mode to operate under as a separate constitutional question. A society can also change Mode later, by the same process.

3.3 The Issuance Channels

All Modes use a common channel architecture of up to five channels — K1, K2, the consumer dividend K3, the inflation-gap stabilizer KI, and the transition-only KT. What differs is which channels are active, how they are calibrated, and whether any adaptive governors are applied.

Channel	Trigger	Calibration	Destination	Lock	Active in
K1	New verified citizen	Fraction of GDP per capita	Stable Floor account	Until age 65	All Modes
K2	Real GDP growth	Fraction of real growth, or $(1 - \kappa_d) \times (g_r \times M2 - K1_{agg})$ in Mode B	Stable Floor account	Until age 65	All Modes
K3	Real GDP growth (shares the real-growth budget)	$\kappa_d \times (g_r \times M2 - K1_{agg})$; within the circulating-pool ceiling	Direct citizen deposit	None — spendable	Any Mode with $\kappa_d > 0$ ($\kappa_d = 0$ in the base Modes)
KI	CPI path gap (when active)	Formula targeting CPI	Direct citizen deposit	None — spendable	Mode C (permanent, inflation-calibrated); Mode Λ

					(conditional); any Mode by constitutional choice — inflation outcome depends on calibrated magnitude
KT	Legacy debt during transition	Price-level path, ~1.5% of M2	Legacy Trust bond redemption	N/A — not to citizens	Mode T only — throttles to dormancy when public debt stabilizes within its operational band

Table 1. The issuance channels. K1 and K2 are common to all Modes. K1 is the direct citizen-deposit channel; whether it produces inflation depends on its calibrated magnitude relative to real economic growth. In the four base Modes K1 is permanently active in Mode C and conditionally active in Mode A, but it can be calibrated at sub-growth magnitudes in other configurations without producing positive inflation. KT is a transition-only channel that funds Legacy Trust debt retirement during the migration described in the transition paper (Neo-Solon, 2026c); it is active only in Mode T, is not a citizen-distribution channel, and throttles to dormancy once the public debt stabilizes within its operational band (the transition paper sets this band at approximately 30 to 60 percent of GDP), thereafter remaining available as a symmetric open-market instrument. K3 is the consumer-dividend channel: like K1 it deposits immediately spendable dollars equally to every citizen, but where K1 is a path-targeting stabilizer triggered by an inflation gap, K3 is growth-indexed — it pays out a share κ_d of the same real-growth budget that funds K2, routed to citizens as current income rather than into locked floors. Because K1, K2, and K3 partition one growth budget rather than stacking, K3 reallocates issuance from the Stable Floor pool to the circulating pool without expanding the money supply above the growth-matched line; κ_d is a Mode parameter (Section 7.5). Adaptive multi-governor Modes may also apply multipliers and boosters to K1 and K2 (see Section 9).

3.4 Two Pools, Four Channels

The dual-circuit architecture is preserved across all Modes. The circulating pool prices the active economy. The Stable Floor pool holds locked equity. K1 and K2 fund the Stable Floor pool through the capital markets channel: when the channel deposits dollars into a Stable Floor account, those dollars are immediately used to purchase shares of the total market index, and the share sellers receive cash that enters normal circulation. The Stable Floor account holds equity, not stored dollars.

K1, when active, funds the circulating pool directly. The dollars are deposited into ordinary citizen accounts and are immediately spendable. This is by design: K1's purpose is to put new money into circulation, not to add to the Stable Floor. The dollars enter at the citizen level — equally distributed, with no intermediating institution — but they enter circulation directly rather than through capital markets.

K3 funds the circulating pool the same way K1 does — equal, unlocked, immediately spendable citizen deposits — but it draws on the growth budget rather than an inflation gap. It is the present-tense counterpart to K2: the two read the same input (real economic growth), and κ_d sets how that year's growth-issuance is split between the Stable Floor pool (K2, future wealth) and the circulating pool (K3, current income).

Because the split is of one fixed budget, turning K3 up turns K2 down by the same amount; the total stays on the growth-matched line, so a positive K3 redistributes between the two pools without itself adding to inflation. It is bounded on the spending side by the same circulating-pool ceiling KI respects, and on the saving side by Stable Floor adequacy — diverting too much growth from floors would erode the age-65 capital guarantee. K3 is the channel that lets a Mode pay a standing, universal dividend out of real output while remaining price-stable.

3.5 Adaptive Smoothing and the Composite Productivity Index

Both K2 and KI use the smoothed Composite Productivity Index applied through the Adaptive Smoothing Rule. Default smoothing is the five-year rolling average. When the most recent year's value deviates from the five-year average by more than two percentage points, the calibrator switches automatically to a two-year rolling average until the deviation returns within one percentage point. The mechanism is rules-based — no committee judgment is required.

Method A (Minimum) is recommended for K2 calibration: use the lower of smoothed GDP and smoothed CPI growth rates, biasing the system toward conservative issuance when measures diverge. KI uses the trailing twelve-month CPI as published by BLS, with the same Adaptive Smoothing Rule applied to prevent over-reaction to single-month spikes.

Three distinct cadences operate in the framework and should not be conflated. (1) The smoothing window is five years by default, switching to two years at turning points — this is input data preparation. (2) The calibration cadence is annual for K2 and KI, and event-based for K1 (recomputed at each new-citizen event). (3) The distribution cadence is monthly for K2 and KI (annual amount divided into twelve equal deposits) and single-event for K1.

3.6 Citizenship and K1

Every person born in the United States is a citizen under the 14th Amendment — K1 deposits at birth are unconditional regardless of parent immigration status. Every person who completes naturalization receives a pro-rated K1 deposit at the moment citizenship is conferred. The pro-ration is based on the fraction of the standard accumulation horizon remaining: a person who naturalizes at age 30 receives 35/65ths of the full K1 endowment (35 years remaining out of the 65-year accumulation window), and a person who naturalizes at age 50 receives 15/65ths. This design is deliberate: it provides a meaningful endowment proportional to the compounding time available while preserving actuarial consistency across the citizen population. A person who naturalizes at age 64 receives a small but nonzero endowment — the framework treats the moment of citizenship as the moment of participation, regardless of age. Refugees and asylees receive K1 only when they later naturalize. Legal residents and unauthorized residents do not trigger K1 issuance, but their economic output contributes to the GDP measure that calibrates K1 deposit sizes for citizens.

In Mode C and Mode A, KI is also distributed only to verified citizens. Non-citizens contribute to the GDP measure that calibrates K1 and K2 and to the CPI that calibrates KI, but do not receive direct distributions. This is a design choice rather than a technical necessity; a future amendment could extend KI to legal residents without altering Model architecture.

3.7 The Stable Floor Account

The Stable Floor is not a retirement account. That distinction is the most important one this section can make, and it is worth stating before the mechanics. A retirement account is a savings vehicle — its purpose is to accumulate funds that replace labor income after work ends. The Stable Floor is a constitutional minimum capital stake. Its purpose is to guarantee that every citizen, regardless of labor-market success, inheritance, education, financial literacy, or behavioral discipline, holds a permanent and growing ownership position in the national productive economy from the moment of citizenship. The dollar value accessible at age 65 is one measurable expression of that stake — but the stake's significance is not confined to age 65.

The existence of a guaranteed non-zero capital floor changes the structure of a citizen's entire life, not just its final chapter. A citizen who knows from birth that absolute destitution in old age is constitutionally impossible operates in a fundamentally different decision environment than one who does not. That difference affects risk tolerance, labor bargaining power, susceptibility to predatory debt, geographic mobility, entrepreneurial behavior, and the psychological relationship between work and survival — throughout working life, not at its end. These are not claims this paper proves. They are the research questions the framework generates, and they are worth naming explicitly because the retirement-income figures that follow, while accurate, do not capture them.

The Stable Floor is also not a public pension or a sovereign wealth fund saving on citizens' behalf. No tax revenue funds it. No government budget line supports it. It is funded entirely by what this framework calls citizen seigniorage — the value created at the moment new money enters the economy, constitutionally routed to citizens as equity rather than captured by an issuing institution. The term seigniorage in its classical sense describes profit captured by a sovereign or institution from currency issuance; no comparable term exists for issuance value that is distributed at creation to the citizens whose economic activity justified creating it. Citizen seigniorage is used throughout this paper to mark that distinction. Under the current system this value flows to financial institutions through the Cantillon hierarchy; under the Citizens Standard it flows into individually owned citizen equity stakes. The citizen is not receiving a transfer. The citizen is receiving a constitutional share of what monetary issuance produces, denominated in productive ownership rather than in promises.

The lock and withdrawal architecture has three components. First, a constitutional lock until age 65: no early access under any circumstances except the bridge-loan provision in Tool 11 (citizens 60+ may borrow up to 10 percent of balance against their own equity at zero interest, repayable from account or estate). Second, a one-time lump-sum

option: at age 65, a citizen may elect to withdraw up to 25 percent of their Stable Floor balance as a single lump sum. Third, an annual withdrawal cap: the citizen may withdraw up to 5 percent of the remaining balance per year, with the unwithdrawn portion continuing to compound at market rates.

This structure is not arbitrary. The 25 percent lump-sum cap is calibrated against international precedent (UK pensions allow a 25 percent tax-free lump sum). The 5 percent annual cap derives from the standard 4 percent safe-withdrawal-rate convention with a small upward adjustment — at 5 percent annual withdrawal on a balance compounding at approximately 4 percent net of inflation in equity markets, the principal preserves real value indefinitely or declines slowly across a typical retirement, and heritability remains meaningful.

The withdrawal cap interacts differently with each Mode. In Mode A, the real Stable Floor at age 65 is approximately \$233,000 at the general-equilibrium realizable return (Neo-Solon, 2026e, Section 6.7); 5 percent annual withdrawal produces approximately \$11,600 per year — but Mode A's principal benefit is not the floor, it is the rising real value of wages and cash under mild deflation. In Mode B, which issues at the full real-growth-matched rate and directs sixty percent of that budget to the floor, the floor is the largest of the base Modes — approximately \$413,000 — and is accompanied by a standing K3 dividend of approximately \$42.75 per month per citizen at launch; 5 percent withdrawal on the floor produces approximately \$20,700 per year. In Mode C, the floor is approximately \$230,000, marginally below Mode A's \$233,000 because, although the two share the same K1 and K2 calibration, Mode A's deflation lifts the real value of its deposits; 5 percent withdrawal produces approximately \$11,500 per year — modest, but meaningful as a complement to the KI dividend.

Heritability deserves explicit framing. When a citizen dies before exhausting their Stable Floor, the remaining balance passes to designated beneficiaries. The inheritor receives the balance into their own Stable Floor account under the same lock-and-withdrawal rules: they may not access it before their own age 65, and at 65 they receive the same 25 percent lump-sum option and 5 percent annual cap. This ensures heritability does not produce younger windfall recipients who can liquidate large sums immediately, while still preserving generational wealth transfer as a structural feature of the framework.

Citizen Stable Floor Accumulation Over a Working Life
(Representative cohort born at framework launch; general-equilibrium realizable return)

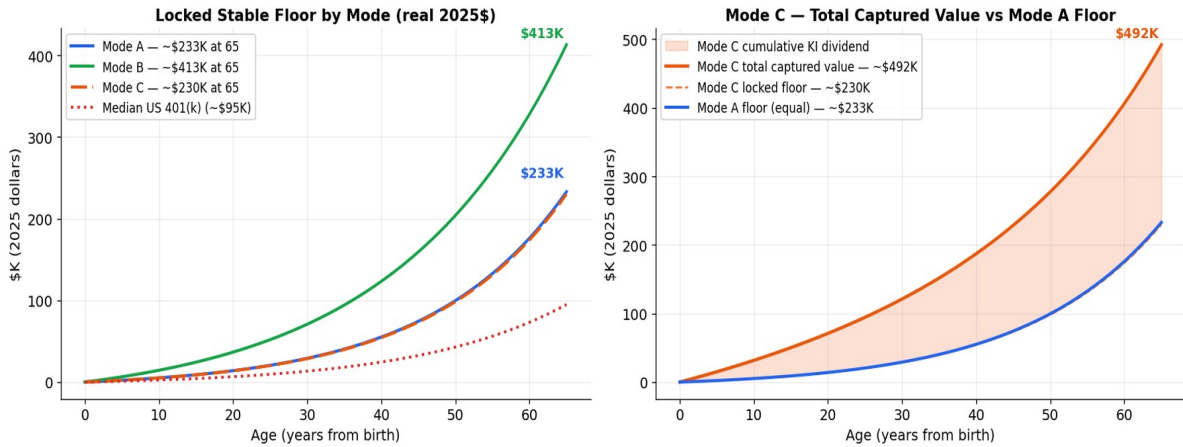


Figure 1. Citizen Stable Floor accumulation over a working life, in real (2025) dollars on the general-equilibrium realizable return. Left: the locked Stable Floor for each base Mode — Mode B at the full real-growth-matched issuance rate, 60 percent of which builds the floor, Modes A and C sharing the same lower-rate floor (~\$230K) — with the current-system median 401(k) (dashed, ~\$95K). Right: Mode C’s total captured value, its Stable Floor plus the launch-year-real value of its cumulative KI dividend receipts (~\$310K real; ~\$142K of that dividend stream in nominal spendable income), against Mode A’s equal floor.

4. Mode A — The Deflationary System

Mode A targets approximately 1.9 percent annual deflation in the circulating pool. Mode A is the appropriate choice for a society that values structurally rising real wages, savings whose purchasing power increases without active management, and the strongest possible separation between the active economy and new money creation.

4.1 Parameters

K1 is calibrated at 2.5 percent of GDP per capita at the moment of citizenship. At launch (US GDP \$30.8T, population 341.8M, GDP per capita \$90,000), K1 produces a deposit of approximately \$2,250 per new citizen. K1 and K2 together draw on a growth-matched issuance envelope equal to 19.5 percent of M2 real growth. K1 is issued to each new citizen on citizenship rather than on growth, and is charged first against the growth budget; K2 is the residual, distributed equally to all living citizens. In a flat or contracting year K1 still flows while the K2 residual falls to zero. At launch (2 percent real growth on \$22.4T M2) the envelope is approximately \$78 billion; net of the \$9 billion K1, K2 produces approximately \$69 billion in annual issuance, distributed to 342 million citizens at about \$203 per citizen in year one. KI is inactive.

Parameter	K1 — Citizenship	K2 — Growth	KI
Trigger	Each new verified citizen	Real GDP growth > 0	Inactive
Calibration	2.5% of GDP per capita	19.5% of M2 growth, less K1	—
Inception value	\$2,250 per new citizen	\$203 per citizen, year 1	—
Annual issuance	\$9B	\$69B	\$0

Destination	Stable Floor (locked)	Stable Floor (locked)	—
Combined as % of M2	0.04%	0.31%	0%

Table 2. Mode A parameters at launch. K1 and K2 calibration uses the smoothed Composite Productivity Index (Section 3.5).

4.2 Outcomes

The circulating pool receives no direct monetary creation. K1 and K2 funds enter Stable Floor accounts through the capital markets channel: the FDCA creates new dollars in its reserves, those dollars purchase total-market index shares from existing equity holders, and the cash flows to share-sellers. The capital-markets channel issues approximately 0.35 percent of M2 per year, but under the dual-circuit architecture that issuance accumulates in the asset circuit; only the bounded spillover — approximately \$16 billion, about 0.14 percent of the transactional circuit M^T (Neo-Solon, 2026e, Section 3.2) — reaches the circulating pool, so against real growth of 2 percent per year the net structural deflation is approximately 1.9 percent. A citizen who holds dollars in a transaction account watches their purchasing power rise without any action on their part.

Under the general-equilibrium realizable return on capital derived in the macroeconomic model (approximately 5.4 percent for the low-capture Modes; Neo-Solon, 2026e, Section 6.7), a citizen born into Mode A reaches age 65 with a real Stable Floor of approximately \$233,000. This is among the smallest floors of the base Modes, by design: Mode A directs the least new money into Stable Floors because its defining benefit is delivered through the circulating pool rather than the locked floor. With the price level falling approximately 1.9 percent per year, every dollar a citizen earns or holds gains purchasing power automatically — over a 65-year horizon, cash and savings appreciate approximately 3.4 times in real terms with no saving or investment decision required. At the 5 percent withdrawal cap the floor yields approximately \$11,600 per year, modest by design, because in Mode A the citizen's gains accrue chiefly through rising real wages and self-appreciating money during working life rather than through the locked stake.

4.3 The Deflation-Debt Question

Mild deflation raises the real burden of existing nominal debt. Mode A resolves this at the architectural level: all debt contracts originated under the Citizens Standard are denominated in price-indexed units. The FDCA publishes a daily price index. Loan principals adjust annually with the index. Chile has operated the Unidad de Fomento system in exactly this configuration since 1967; Chilean financial markets function normally with this provision.

4.4 When Mode A Is Appropriate

Mode A is appropriate for societies that value purchasing power preservation as a primary monetary outcome, that have legal infrastructure capable of administering universal price-indexed debt contracts, and that accept the wage-rigidity costs that deflation can impose on workers without bargaining power to negotiate indexed wages.

Mode A is the strongest possible response to the 96 percent purchasing power loss the current system has produced over the past century.

4.5 The Hoarding Question and the Demurrage Backstop

A deflationary regime invites one standing objection: if holding money earns a positive real return simply by sitting still, will households hoard it, and will velocity collapse into a self-feeding contraction? The empirical record answers the general version of the question first. Across seventeen countries and more than a century, deflation and depression are linked only in the 1930s; outside that single episode there is virtually no association between falling prices and falling output (Atkeson and Kehoe 2004), and once asset prices are controlled for, the goods-price decline carries no depression signal even in the 1930s (Bordo, Landon-Lane and Redish 2004). The deflation Mode A produces is the benign, productivity-driven kind, prices falling because output per hour rises, not the demand-collapse kind. On the central money-demand elasticity (roughly 0.3 to 0.6 across long international samples; Lucas 2000; Benati, Lucas, Nicolini and Weber 2021), Mode A's approximately 1.9 percent deflation depresses velocity by only about one percent relative to a price-stable baseline. The steady state needs no intervention.

The hoarding risk lives not in mild realized deflation but in the tail where deflation expectations un-anchor or a confidence shock triggers precautionary hoarding, the dynamic that drove the velocity collapse of the early 1930s, which the historical literature attributes to a surge in risk premia and uncertainty rather than to the price decline itself. For that tail the framework needs a tool the inflation channel KI cannot supply: KI fights below-target prices by issuing money, which in a deflation-targeting regime would destroy the very price path Mode A is chosen to deliver. The instrument that fits is a demurrage backstop, a small carrying charge on transactional balances that offsets the real return to idle cash, decoupling the price path from the incentive to hoard. It is the one lever that addresses hoarding without abandoning the deflation target.

The backstop is conditional, ring-fenced, and revenue-neutral by construction. It is dormant by default and arms only when velocity falls materially below baseline, the signal that benign deflation is tipping into hoarding; intervening in the benign state would impose a deadweight carrying cost for no benefit. When armed, the charge is set near the deflation rate itself, on the order of two percent, far below the eight to twelve percent rates of the historical demurrage scrips (Worgl 1932; the Chiemgauer), and it applies only to the transactional circuit, never to the locked Stable Floors, the same ring-fence the reverse-KT drain respects. Its proceeds are redistributed equally per capita, so no money is created or destroyed and the price path is untouched; only the relative incentive moves, with above-average idle balances paying in and fast-circulating balances receiving back. The charge changes behavior at the margin without functioning as a net tax.

Two honest limits bound the tool. First, demurrage can be evaded by shifting into untaxed liquid stores of value, the objection Keynes raised against Gesell's original

proposal; at national scale, with assets and foreign currency available as escapes, the usable charge is capped near the liquidity premium households pay to hold transactional money, on the order of two to four percent. The historical scrips sustained higher rates only because they were the sole local means of payment and were receivable for taxes, a captive demand a national currency lacks. Second, that ceiling means the charge can fully offset the modest hoarding incentive of mild deflation but only partially counter a deep expectations spiral before flight dominates; in the severe tail it is a damper, not a clamp, and would be deployed alongside other instruments rather than alone. The backstop should therefore be understood as one bounded, rules-triggered, automatically reversing entry in the emergency toolkit of Section 12, listed there as Tool 16, specific to Mode A and any other deflation-targeting configuration, and not as a standing feature of the core architecture. The supporting analysis, calibration against the money-demand and demurrage-experiment evidence, and the velocity model are released as the demurrage modules in the replication package.

5. Mode B — The Split-Dividend System

Mode B issues at the full real-growth-matched rate on a fixed 60/40 split. For an economy whose transaction-active money share sits near one-half it produces approximately zero annual drift; away from that structure the fixed split drifts (Section 5.x), and guaranteed price stability in any monetary structure is the role of Mode Ω (Section 8). K1 calibration matches Mode A — a fixed share of GDP per capita per new verified citizen. K1 and K2 together fund the circulating pool at the full real-growth-matched rate — K2 issued as the residual after K1 — so that new money enters at the pace of real output and the price level holds stable by construction. This holds under steady growth: because K1 fires on new citizenship rather than on growth, total issuance is the growth-matched envelope plus a guaranteed citizenship floor of roughly 0.04 percent of M2, so a flat-growth year carries about +0.04 percent rather than exactly zero, and Mode A's deflation — which requires real growth to outpace issuance — attenuates toward zero without it. Mode B preserves the dual-circuit architecture and citizen-anchored issuance; it holds the price level approximately stable for economies near a one-half transaction ratio (Mode Ω , Section 8, targets stability in any structure). Of the post-K1 growth-matched budget, Mode B directs sixty percent into the locked Stable Floor and pays forty percent as a standing, spendable K3 dividend; because both draw on the one budget rather than stacking, the split holds the price level exactly stable (Neo-Solon, 2026e, Sections 4.3 and 6.7). This is the configuration in which the system lands after the debt-retirement transition (Mode T-stable, Section 10) for an economy at the balance point; away from it, Mode T-stable is the Mode Ω solve of Section 8.

5.1 Parameters

Total annual issuance is calibrated to the full real-growth-matched amount. At launch (2 percent real growth on \$22.4T M2), this produces approximately \$447 billion in annual issuance, split between K1, K2 (floor), and K3 (dividend) so the channels together sit on the line. K1 is calibrated identically to Mode A — 2.5 percent of GDP per capita per new

verified citizen — producing approximately \$2,250 per new citizen (~\$9 billion aggregate). K2 takes 60 percent of the residual after K1: approximately \$263 billion per year, or approximately \$769 per existing citizen per year, with the remaining 40 percent (~\$175 billion) paid through the K3 dividend channel. K1 remains inactive.

Parameter	K1 — Citizenship	K2 — Growth	K3 — Dividend
Trigger	Each new verified citizen	Real GDP growth	Real GDP growth
Calibration	2.5% of GDP per capita	60% of real-growth-matched (real growth × M2) less K1	40% of real-growth-matched
Inception value	\$2,250 per new citizen	~\$769 per existing citizen, year 1	~\$513 per existing citizen, year 1
Annual issuance	~\$9B	~\$263B	~\$175B
Destination	Stable Floor (locked)	Stable Floor (locked)	Citizen dividend (spendable)
Combined as % of M2	0.04%	1.17%	0.78%

Table 3. Mode B parameters at launch. Total issuance is approximately 2 percent of M2 — the full real-growth-matched amount — producing approximately stable prices near a one-half transaction ratio (approximately zero annual drift).

5.2 Outcomes

The price level remains stable, with approximately zero annual drift. Nominal wages and nominal contracts retain their meaning across years and decades. Long-term planning under Mode B is straightforward: a dollar saved today is approximately a dollar at retirement, and real returns on investment come primarily from real economic activity rather than from the monetary regime.

Evaluated against actual US historical data in the companion empirical paper (Neo-Solon, 2026b), Mode B produces a Stable Floor that lands at approximately median actual retirement wealth across four cohorts spanning 1960–2055 — below the median for the earliest cohorts and above it for the latest — documented on the same general-equilibrium realizable basis used here, and counting the standing K3 dividend the citizen also receives the total value captured clears the median for every cohort. Because Mode B holds the price level approximately stable near a one-half transaction ratio, nominal and real values coincide there. The empirical paper documents the full cohort range, the forward transition-cohort projection, and the Monte Carlo distribution; this paper defers all empirical magnitudes to that analysis. Calibration note. Mode B calibrates K1 and K2 together at the full real-growth-matched rate — K2 as the residual after K1 — so that the money supply grows at the pace of real output and the price level holds stable by construction in growth years (zero drift, nominal equal to real); in the rare contraction years, when the growth-matched K2 is zero, the citizenship deposit K1 continues to flow and acts as a small automatic stabilizer. This is the calibration the system settles into after the debt-retirement transition of Section 10 for an economy at the balance point; there Mode T-stable coincides with the fixed Mode B split. Away from the balance point Mode T-stable is the Mode Ω solve (Section 8), which holds the price level for any transaction ratio; the United States, near the balance point, is the case where the two agree. The companion empirical paper (Neo-Solon, 2026b) evaluates this calibration

against historical US data on the same realizable basis. Mode A's adaptive governors, described in Section 9, represent a further refinement: rather than a static calibration, Mode A adjusts K2 dynamically in response to observable demographic and productivity conditions.

5.3 When Mode B Is Appropriate

Mode B is appropriate for societies that value nominal contract stability above purchasing power gain, that have extensive long-duration financial contracts (mortgages, pensions, insurance) whose pricing assumes stable nominal terms, and that prefer the reduced wage-rigidity friction that stable prices produce.

A society in Mode A can transition to Mode B in either of two ways. First, by Tier 2 constitutional amendment with the standard 67 percent supermajority. Second, by the Long-Run Stability Provision: when the GDP-to-M2 ratio crosses a constitutional threshold, parameters recalibrate automatically to functional Mode B equivalence and recalibrate back if the ratio later falls. Under base-case 2 percent real growth, the threshold is reached approximately 80 to 120 years into Mode A operation. Both pathways are available; the cycling mechanism is described in detail in Section 18.1.




Figure 2. Purchasing power of \$1 over 65 years under each Mode and the current system. Mode A produces structural purchasing power gains through mild deflation. The current system destroys approximately 85% of purchasing power over 65 years at 3%/yr inflation.

6. Mode C — The Inflationary System

Mode C targets approximately 2 percent annual inflation, scaled to track real growth and population growth. Mode C activates the inflation channel — KI — that distributes new money equally to all citizens as an unlocked, immediately spendable citizen dividend. KI

calibration is annual; KI distribution is monthly. K1 and K2 continue to operate exactly as in Mode A, preserving the Stable Floor.

Mode C is the configuration that makes the Citizens Standard's citizen seigniorage explicitly visible to citizens. In Mode A and Mode B, citizen seigniorage funds Stable Floor accounts that citizens cannot access until age 65. In Mode C, seigniorage is split: K1 and K2 still fund the Stable Floor, but KI returns a meaningful portion of new money creation to citizens as current income.

6.1 The KI Formula

KI annual issuance is calibrated by a price-level path targeting formula. The framework defines a desired CPI trajectory growing at the target inflation rate from launch, and KI closes the cumulative gap between the actual CPI and that trajectory each year. This is structurally different from a simple single-year inflation gap formula, which oscillates around the target. Price-level path targeting self-corrects: undershooting one year produces a larger gap to close the next year, and overshooting reduces the gap and contracts KI accordingly.

$$\begin{aligned} \text{target_price_level}(t) &= (1 + \text{target})^t \\ \text{gap}(t) &= \ln(\text{target_price_level}(t-1) / \text{actual_price_level}(t-1)) \\ \text{KI}(t) &= \max(0, [\text{target} + \text{real_growth} + \lambda \times \text{gap}(t)] \times \text{M2} - \text{K1_agg} - \text{K2_agg}) \end{aligned}$$

Here *target* is the constitutionally specified inflation target (default 2 percent, written π^* below); *real_growth* is total real GDP growth, taken from the smoothed Composite Productivity Index. Population growth is tracked separately and does not enter the aggregate issuance rate; it appears only when the aggregate is divided to a per-citizen dividend. λ is a closure-speed coefficient (default 0.5); *actual_price_level* is the prior year's CPI as published by BLS and *target_price_level* is the constitutional path $(1 + \text{target})^t$; *M2* is the prevailing circulating money supply; and *K1_agg* and *K2_agg* are the aggregate K1 and K2 budgets, netted out because they already issue at or below the growth line. The formula separates two jobs. The maintenance term (*target* + *real_growth*) sets the money-growth rate that sustains the target path: by the quantity identity, money growth of $(\pi^* + \text{real_growth})$ against real growth of *real_growth* yields exactly π^* inflation. The error-correction term ($\lambda \times \text{gap}$) repairs deviations from the path — lifting issuance when the price level runs below path, cutting it (and flooring at zero) when above — and because *gap* is the log distance to a path the rule itself drives to zero, the correction term vanishes on the path and does not bias the steady state. The abbreviation *ln* denotes the natural logarithm; the log form makes the gap additive across years, so a multi-year cumulative deviation is closed at a consistent rate regardless of its size.

KI is best read as a citizen-distributed price-level-path-targeting rule, and it inherits a long line of monetary thought. Its commitment to a price-level path rather than a period-by-period inflation rate is Wicksellian (Wicksell, 1898) and is the modern price-level-targeting prescription of Woodford (2003) and Eggertsson and Woodford (2003), under

which past misses are not treated as bygones and the price level is returned to its path. Its insistence on a transparent, legislated rule over period-by-period discretion follows Friedman (1960), and its use of a feedback term on a nominal target — rather than a fixed money-growth constant — follows McCallum's (1988) monetary-base feedback rule. The older idea that money should be issued to defend a stable purchasing path traces to Fisher's compensated dollar (Fisher, 1920). What is new in KI is neither the target nor the feedback but the distribution: the issuance required to defend the path is paid out as an equal per-citizen dividend, making KI a synthesis of price-level-path targeting with a citizen's share rather than an open-market operation.

6.2 Parameters at Launch

On a clean launch the price level begins on its own target path, so $gap(t) = 0$ and KI issues at its maintenance rate from the first year: approximately 1.98 percent of M2 (3.85% of M^T), or about \$442 billion, distributed to approximately 342 million citizens at roughly \$108 per citizen per month (\$1,293 per year). The rate is flat — there is no inflation ramp, and Mode C runs at its 2 percent target from year one. Because KI is a constant share of a transactional circuit that grows with money and real output, the dollar amount scales gently with the economy: it passes roughly \$471 billion, or about \$115 per citizen per month, within about two years, and continues to rise thereafter while the rate holds near 1.98 percent of M2. K1 and K2 contribute approximately \$80 billion combined through capital markets in year one. Should the Standard instead launch into an economy whose price level is off the new target path, the error-correction term closes that gap over the following few years before issuance settles to the maintenance rate.

Parameter	K1 — Citizenship	K2 — Growth	KI — Inflation-Gap Stabilizer
Trigger	Each new verified citizen	Real GDP growth	Cumulative CPI gap, computed annually
Calibration	2.5% of GDP per capita	19.5% of M2 growth, less K1	Price-level path targeting
Launch year value	\$2,250 per new citizen	\$203 per citizen	\$108/mo per citizen (\$1,293/yr)
Steady-state (year 5+)	Scales with GDP/cap	Scales with growth	~\$115/mo within ~2 yrs (\$1,378/yr)
Annual issuance (launch)	\$9B	\$69B	~\$442B
Destination	Stable Floor (locked)	Stable Floor (locked)	Spendable citizen transaction account
Combined as % of M2	0.04%	0.31%	~1.98%

Table 4. Mode C parameters under price-level path targeting. KI calibration is annual; KI distribution is monthly.

6.3 How New Money Enters the System in Mode C

Mode C is the only configuration in the framework where new money enters circulation through two structurally different transmission channels operating simultaneously. K1 and K2 enter through the capital markets channel, identical to their operation in Mode A and Mode B — new dollars purchase equity shares, and cash flows to share-sellers. KI enters through the direct deposit channel: the FDCA creates new dollars and deposits

them, in equal per-citizen amounts, directly into ordinary citizen transaction accounts, spendable that same day.

The distinction matters for inflation. Capital-market injection produces minimal direct consumer-price inflation pressure because new money is mediated by equity markets before reaching consumers. Direct-deposit injection produces inflation pressure roughly proportional to its magnitude — which is why Mode C, and not Mode A or Mode B, hits a positive inflation target. A society choosing between Modes is choosing whether to activate the direct-deposit channel.

Why The Distinction Matters

A reader could object that KI reproduces the discretionary system's pattern: new money entering the consumer economy directly. The objection captures something real, but misses the architectural point. The discretionary system's direct injection occurs through bank lending, where new money enters the accounts of borrowers who tend to be wealthier than median. KI's direct injection enters every citizen's account in equal amounts on the same day. The transmission mechanism is the same; the distributional structure is opposite.

Equally important: KI's magnitude is constitutionally bounded by the price-level path targeting formula, not by committee judgment. When inflation runs above the target trajectory, KI contracts automatically — not because the FDCA decides to slow issuance, but because the formula's gap-closure term turns negative. The discretionary system has no equivalent self-correcting mechanism.

6.4 The KI Lock — and Why There Isn't One

KI is unlocked by design. KI's purpose is to put new money into the circulating pool to hit an inflation target — a locked KI would not enter circulation. The Stable Floor is provided by K1 and K2, which remain locked exactly as in Mode A. KI is additional — a current-income complement that runs throughout working life, not a substitute for the Stable Floor mechanism. The Cantillon Effect does not operate when distribution is equal at issuance: every citizen receives the same KI deposit on the same day, with no late recipients.

6.5 Outcomes

On a clean launch the price level begins on its target path, so the maintenance term issues at the full rate and inflation runs at approximately the 2 percent target from the first year. Should the Standard instead launch with the price level below the new target path, Mode C exhibits a brief convergence period — actual inflation running below target at first and rising to the 2 percent target over the first several years as the error-correction term closes the gap — after which the price level tracks the constitutionally specified path with year-to-year deviation typically below 0.5 percentage points.

Citizens receive a meaningful monthly dividend that scales with the economy. At launch, KI produces approximately \$108 per citizen per month or \$1,293 per year, at a flat

issuance rate of about 1.98 percent of M2 (3.85% of the transactional circuit $M^{\wedge}T$). Because that rate is a constant share of a growing transactional circuit, the dollar dividend scales gently with the economy — passing roughly \$115 per citizen per month (about \$1,378 per year) within about two years and rising thereafter. A family of four receives approximately \$432 per month at launch, scaling past roughly \$460 per month within a couple of years. Cumulative KI receipts across a 65-year working life total approximately \$142,000 in nominal dividends, spent as current income across the span; deflated to launch-year purchasing power at Mode C's 2 percent inflation, that stream is worth approximately \$80,000 in real terms.

Mode C shares Mode A's K1 and K2 calibration, so the two build nearly the same Stable Floor — approximately \$230,000 in real launch-year purchasing power at age 65 (marginally below Mode A's \$233,000, since deflation lifts the real value of Mode A's deposits), earning the same general-equilibrium return because they capture the same share of the capital stock (Neo-Solon, 2026e, Section 6.7). The Stable Floor holds total-market equity, whose real value compounds at that return regardless of the price-level regime, so Mode C's inflation does not erode the floor itself. What distinguishes Mode C is not the size of its locked stake but the KI dividend it pays alongside it: an unlocked, immediately spendable citizen dividend of approximately \$108 per month at launch at a flat rate of about 1.98 percent of M2, scaling past roughly \$115 per month within about two years and rising with the economy thereafter — to approximately \$276 per month by the end of the 65-year horizon. Mode C trades Mode A's deflationary real-wage gains for visible current income while building the same modest floor.

Stating the lifetime tradeoff honestly: a citizen under Mode C accumulates approximately \$310,000 in total real captured value (a Stable Floor of approximately \$230,000 plus a KI dividend stream worth approximately \$80,000 in launch-year real terms), against a Mode A citizen's Stable Floor of approximately \$233,000. In nominal terms the KI stream is larger, approximately \$142,000 of spendable income across the working life, since it is paid and spent as current dollars rather than held; the real figure is the honest one for a lifetime comparison, while the nominal figure is what reaches the citizen's account month to month. The two Modes build essentially the same floor; Mode C adds spendable dividend income on top of it. Mode A's distinct benefit does not appear in this stock comparison: it is the rising real value of wages and cash under mild deflation, which lifts purchasing power across the whole working life rather than accumulating in any single account.

6.6 Why This Is Not Universal Basic Income

Mode C's KI superficially resembles a Universal Basic Income. The differences are structural. UBI is funded through taxation: someone's income is reduced by tax to fund someone else's payment. KI is funded through citizen seigniorage: the value created when new money enters the economy is redirected to equal citizen distribution rather than to bank balance sheets. In Mode C that issuance is calibrated to produce modest inflation, but the seigniorage mechanism itself is independent of the inflation target — KI could be calibrated at a lower magnitude and still deliver a citizen dividend without

crossing into positive inflation territory. No citizen's income is reduced to fund another citizen's KI. UBI proposals have a fixed nominal amount requiring periodic legislative adjustment; KI is a fraction of M2 calibrated to the inflation target, structurally preserving its real value without legislative intervention. UBI is reversible by ordinary legislative process; KI is a Tier 2 constitutional provision requiring 67 percent supermajority to modify.

6.7 The Cantillon Argument in Mode C

Mode C creates substantially more circulating-pool money than Mode A — approximately 2.3 percent of M2 annually rather than 0.35 percent. A serious reader will ask whether Mode C reproduces the Cantillon problem at higher magnitude. It does not, but the reason requires precision. The Cantillon Effect operates when distribution is hierarchical. Mode C's KI has neither temporal nor distributional asymmetry: every citizen receives the same amount on the same day. There are no early recipients to capture the inflationary bonus and no late recipients to absorb its cost.

6.8 When Mode C Is Appropriate

Mode C is appropriate for societies that explicitly want a citizen dividend funded by citizen seigniorage rather than by taxation, that prefer modest predictable inflation to deflation or strict price stability, and that want monetary policy to produce visible material benefit for citizens rather than merely abstract macroeconomic stability.

Mode C also produces a useful structural feature absent in the other base Modes: a counter-cyclical transmission mechanism that does not require committee judgment. When inflation runs above target, KI contracts automatically. When inflation runs below target, KI expands automatically. Mode C is the only base Mode in which the calibration formula itself provides automatic stabilization without invoking the emergency toolkit.

7. Cross-Mode Comparison

The base Modes are coherent applications of one Model. Choosing among them — or ratifying an extended configuration like Mode Λ — is a constitutional decision that reflects a society's values and tradeoffs. This section presents the comparison directly.

7.1 At Launch

Metric	Mode A — Deflationary	Mode B — Split Dividend	Mode C — Inflationary	Mode D — Distributed
Inflation target	-1.86%	0% (drift)	+2%	0% (drift)
Annual K1 issuance	\$9B	~\$9B	\$9B	\$0
Annual K2 issuance	\$69B	~\$263B	\$69B	\$0
Annual KI / K3 issuance	\$0	~\$175B (K3)	~\$442B launch, ~\$471B by ~yr 2	~\$230B (K3)
Total annual issuance	\$80B (0.35% M2)	\$447B (2.0% M2)	~\$521B (2.3% M2)	~\$230B (1.03% M2)
KI / K3 dividend per citizen	—	~\$42.75/mo launch (K3)	~\$108/mo launch → ~\$115/mo by	~\$56/mo (K3)

			~yr 2 (flat ~1.98%)	
Stable Floor at 65 (real, 2025 \$)	~\$233K	~\$413K	~\$230K	\$0 (no floor)
Cumulative dividend (real, 2025 \$)	—	~\$57K (K3)	~\$80K	~\$75K (K3)
Annual real income at 65 (5%)	~\$13.6K/yr	~\$20.7K/yr	~\$13.5K/yr + KI during working life	\$0 floor; ~\$672/yr dividend
Real wage trajectory	Rising	Flat	Slight erosion	Flat
Nominal contract stability	Requires indexing	Native	Native	Native

Table 5. Cross-Mode comparison at launch parameters and 65-year horizon under base-case assumptions (general-equilibrium realizable return by Mode, Neo-Solon 2026e §6.7; 2% real growth, 0.5% population growth). All projected Stable Floor values assume framework launch at 2025 parameters with 65 years of forward compounding; the companion empirical paper (Neo-Solon, 2026b) documents historically calibrated outcomes for four cohorts (born 1960–1990) that are substantially lower in absolute magnitude than these 2025 launch projections, reflecting the smaller economic base at historical deposit dates, and documenting on the same realizable basis that the floor universalizes approximately the median across the four cohorts — near or modestly below it for the earliest cohorts, above it for the latest, and above it for all once the standing K3 dividend is counted. Policymakers evaluating implementation in non-US contexts should calibrate expectations to their country's long-run equity return history — the twentieth-century median real capital appreciation across global markets was approximately 1.5 percent annually versus approximately 4.73 percent for the US (Goetzmann–Jorion 1997) The companion transition paper (Neo-Solon, 2026c) specifies the migration path from the current system to that destination.

Mode D, shown as the fourth column above, shares Mode B's issuance regime, differing only in distribution. Against Mode B at launch, Mode D pays a larger standing dividend (approximately \$672 per year versus approximately \$513) and builds no Stable Floor (\$0 versus approximately \$413,000), trading the universal locked stake for maximal current liquidity and individual portfolio choice. Because Mode D deepens no aggregate capital stock, the realizable return a citizen earns on the privately reinvested dividend is the unattenuated 6.67 percent (versus 4.26 percent under Mode B); its per-Mode figures are tabulated in Neo-Solon, 2026e, Section 6.7.

To put these numbers in context, the table below compares each base Mode against current US monetary and retirement system outcomes. Two methodological notes: first, the Citizens Standard simulations are deterministic baselines under specified assumptions; current-system numbers are empirical medians from US data sources. Second, Social Security continues to exist alongside the Citizens Standard under all Modes; the framework supplements rather than replaces it.

Table 6. Sixty-five-year outcomes by base Mode versus the current system, on the general-equilibrium realizable basis.

Metric (65-year horizon)	Current System	Mode A	Mode B	Mode C	Mode D
Annual inflation	~3% (1990-2024 avg)	-1.86%	0% (drift)	+2%	0% (drift)
Cumulative purchasing power loss	~85% over 65 yrs	None (gain)	Mild gain	~71%	None (stable)

Median retirement balance at 65	~\$95K (Vanguard 2025)	—	—	—	—
Stable Floor balance at 65 (real)	—	~\$233K	~\$413K	~\$230K	\$0 (none)
Stable Floor income/yr at 5%	—	~\$13.6K	~\$20.7K	~\$13.5K	—
Citizen dividend (monthly)	None	None	~\$42.75/mo (K3)	~\$108-115+	~\$56/mo (K3)
Cumulative dividend receipts (real)	—	—	~\$57K (K3)	~\$80K	~\$75K (K3)
Annual Social Security (average)	~\$25K	~\$25K	~\$25K	~\$25K	~\$25K
Total real retirement income/yr	~\$30K	~\$37K	~\$46K	~\$37K + ongoing KI	~\$26K (SS+div)
Equality of outcome	Highly skewed	Uniform	Uniform	Uniform	Uniform

Current-system figures: the median 401(k) balance (Vanguard, How America Saves 2025) and the average Social Security retired-worker benefit (\$2,079/month, \$24,953/year; SSA Monthly Statistical Snapshot, March 2026). The median is used for 401(k) balances, which are right-skewed so the mean overstates the typical holder; the mean is used for Social Security, whose progressive formula and benefit cap compress the distribution so mean and median nearly coincide and only the mean is published. The median 401(k) counts only account holders, so it is generous to the status quo. Citizens Standard outcomes are uniform per-citizen by architectural design.

Four Base Modes vs. Current System — 65-Year Trajectories

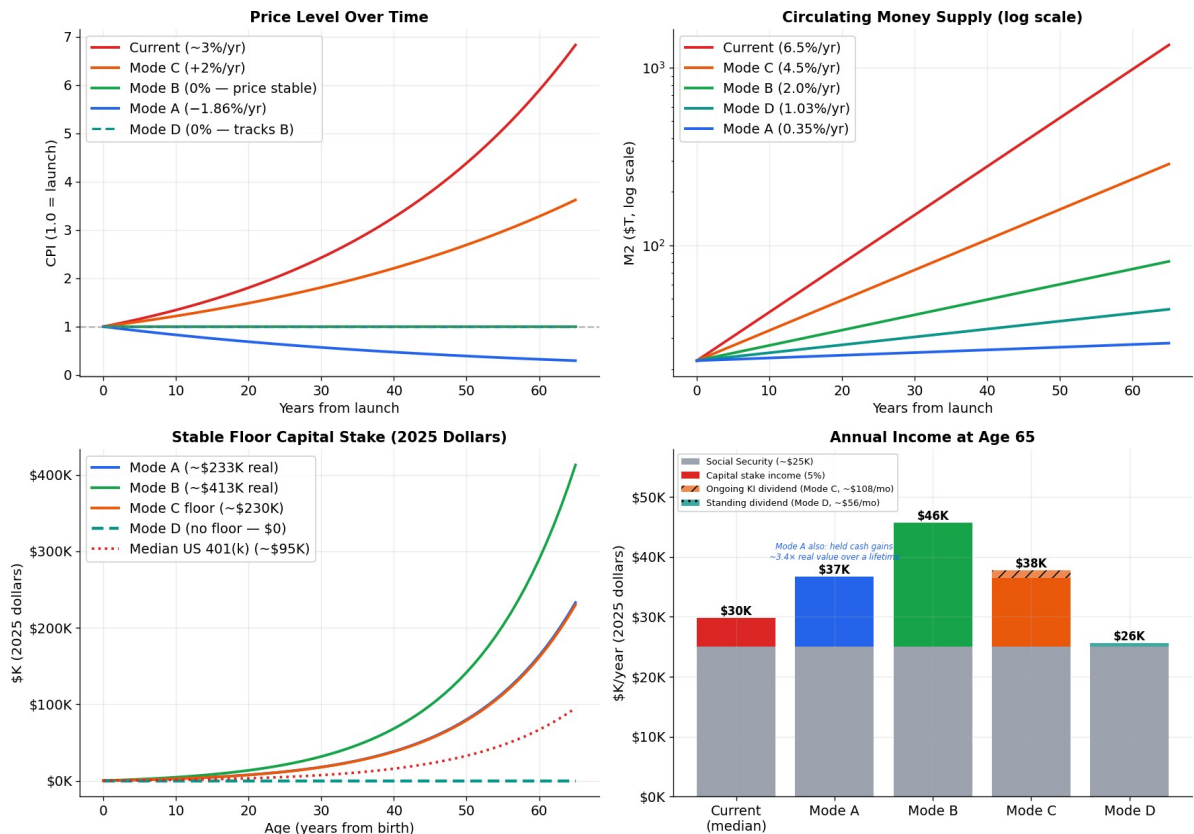


Figure 3. The four base Modes vs. current system — 65-year trajectories. Top left: price level. Top right: circulating money supply (log scale). Bottom left: Stable Floor capital stake in 2025 dollars, with current-

system median 401(k) for comparison (Vanguard 2025 How America Saves, ~\$95K at age 65). Bottom right: annual income at age 65 from the capital stake (5% withdrawal) plus Social Security, with Mode A's deflationary purchasing-power gain (~1.86%/yr real), Mode C's ongoing K1 dividend (~\$108/mo), and Mode D's standing dividend (~\$56/mo) shown as the same citizen benefit delivered several ways. Modes A, B, and C produce substantially larger capital stakes than the current-system median, while Mode D builds no stake by design, delivering its benefit entirely as the standing dividend. Base assumptions: general-equilibrium realizable return by Mode (Neo-Solon, 2026e \$6.7), 2% real growth, 0.5% population growth, smooth baseline.

7.2 Tradeoff Summary

Mode A maximizes purchasing power preservation. Citizens who hold cash, earn wages, or save outside the Stable Floor benefit structurally from monetary deflation, and Mode A is the only base Mode that delivers this structural cash-holder benefit. Its cost is the friction deflation imposes on nominal contracts (resolved through universal price indexing) and on workers without bargaining power to negotiate indexed wages.

Mode B targets approximate price stability, producing the largest Stable Floor capital stake of the base Modes — nearly twice the floor under Mode A or Mode C in real terms at age 65, because it issues K2 at the full real-growth-matched rate rather than the lower deflationary-mode rate, and it pays a standing K3 dividend alongside the floor. The dollar holds approximately its value over time, and the conventional structure of long-duration financial contracts works without modification. The cost is that purchasing power gains do not occur structurally; cash holdings do not appreciate.

Mode C maximizes visible citizen benefit from monetary creation throughout working life. Every citizen receives a meaningful monthly dividend funded by citizen seigniorage rather than taxation, on top of a Stable Floor equal to Mode A's. The structural tradeoff is borne against Mode B, not Mode A: by running K1 and K2 at the low deflationary-rate calibration rather than the full real-growth-matched rate, Mode C builds a smaller locked floor — approximately \$230,000 versus Mode B's approximately \$413,000 — redirecting that issuance into the current K1 dividend. The choice is between a large locked stake at 65 and visible spendable income throughout working life.

7.3 Choosing a Mode

The framework does not recommend a Mode. The choice is properly the citizens', made through ratification or amendment. A reasonable society might choose differently than another reasonable society, and might also change Mode over time as conditions and preferences change.

The framework's claim that Mode B universalizes approximately the median citizen's capital position is evaluated empirically for four cohorts spanning 1960–2055 in the companion empirical paper (Neo-Solon, 2026b). On the general-equilibrium realizable basis used here, the locked Stable Floor lands at roughly median actual 401(k) wealth — modestly below it for the earliest cohorts, above it for the latest — and at only about four-tenths of the mean, the egalitarian signature of a universal floor that delivers the median without reproducing the concentration embedded in the average. Counting the standing K3 dividend the citizen also receives lifts the total value captured above the median for

every cohort. With the qualification that Stagflation-magnitude sequences during peak accumulation push the floor below the median for the most adversely timed cohorts, the empirical paper reports that under bootstrap resampling on the same realizable basis the typical (P50) floor runs from just below the median for the earliest cohort to well above it for the later cohorts; a substantial minority-to-majority of paths fall below the median for the floor alone; and reinvestment of the liquid dividend lifts the median outcome above the benchmark for every cohort. The P5 adverse-tail outcome falls below median actual in every configuration. The full deterministic, P50, and adverse-tail magnitudes are reported there.

The Citizens Standard's claim is narrower and more defensible: regardless of which Mode a society chooses, the Model produces a more coherent, more accountable, and more durable monetary system than the discretionary alternative. The choice among Modes is a debate the framework makes possible. The choice between the Model and committee discretion is the debate the framework wants to win.

7.4 Mode D — The Distributed System

Mode D is the framework's pure-dividend corner — the mirror image of Mode A. Where Mode A routes the entire growth-matched budget into locked Stable Floor accounts and pays no spendable stream, Mode D inactivates the citizenship deposit ($K1 = 0$) and pays the whole budget out as a spendable citizen dividend ($\kappa_d = 100\%$), building no Stable Floor at all. It shares the same issuance regime as Mode B — approximately stable near a one-half transaction ratio, approximately zero annual drift — but distributes new money in the opposite form: entirely as current income rather than as a locked capital stake. The two price-stable Modes therefore differ only along the distribution axis, Mode B holding the larger share of issuance as universal capital and the remainder as dividend, Mode D holding none as capital and all as dividend.

At launch the growth-matched dividend is calibrated to the transactional money-growth need of the economy, so that the amount paid equals $g_r \cdot M^T$ — approximately \$230 billion per year, or about 1.03 percent of $M2$. With $K1$ inactive, the inflation-gap channel KI inactive, and no floor accumulation, this single channel carries the entire issuance: a uniform dividend of approximately \$672 per citizen per year, roughly \$56 per month. There is no birth endowment and no locked account; every dollar created is paid directly and equally to citizens as spendable income.

Mode D is the configuration that sits exactly on the price-stability leash. Stability requires that the money reaching the transactional circuit not exceed that circuit's growth need, $g_r \cdot M^T$; in the other Modes a portion of issuance accumulates in the asset circuit and only the bounded spillover (Neo-Solon, 2026e, Section 3.2) crosses into the circulating pool. Mode D funds no asset circuit — no Stable Floor is built, so no issuance accumulates there and the spillover vanishes — which leaves the dividend itself as the only injection. Because that dividend is sized to $g_r \cdot M^T$, the total lands precisely on the leash rather than below it: the price level holds flat not because issuance is withheld but because the dividend reallocates the full growth budget into the transactional circuit at exactly the rate that circuit grows. Mode D is thus the maximal dividend the framework

can pay while holding the price level neutral; paying more would push the injection above the leash and generate inflation.

Because no Stable Floor is funded, the universal capital stake is zero ($\psi^* = 0$). The framework hands Mode D citizens spendable income and leaves the saving decision to them: a citizen may consume the dividend or reinvest it privately and assemble a capital position of their own choosing. Mode D substitutes individual portfolio choice for the universal locked stake that Modes A, B, and C build by construction — the floor is not abolished but made opt-in, available to any citizen who elects to save rather than mandated for all.

Forgoing a universal floor changes the return a saver faces. In the floor-building Modes, universal accumulation bids up asset prices and compresses the realizable return — under Mode B the general-equilibrium realizable return settles near 4.26 percent once the universal stake is accounted for. Mode D deepens no aggregate capital stock, so this attenuation does not occur, and a citizen who reinvests the dividend faces the unattenuated no-program return on capital, approximately 6.67 percent (Neo-Solon, 2026e, Section 6.7). A Mode D saver therefore earns a higher rate on privately held capital than a Mode B citizen's locked floor accrues: the framework trades the certainty of a universal stake for a higher return on a voluntary one.

The cost of that trade is dispersion. Because retirement security under Mode D depends on individual saving rather than on a guaranteed account, outcomes vary across citizens in precisely the way the Stable Floor was designed to prevent: a citizen who saves little reaches retirement on Social Security and the ongoing dividend alone — on the order of \$23 thousand per year in combined real income — while a disciplined saver may exceed the floor Mode B would have guaranteed. Mode D suits a society that weighs current liquidity and individual choice above a universal capital guarantee and is prepared to accept the distributional variance this reintroduces. It is the cleanest expression of the framework's no-index, all-dividend path: a constitutional, formula-bounded, equally distributed money supply that gives every citizen the same spendable stream and leaves the rest to them.

7.5 Beyond the Four Base Modes

The four base Modes presented in this paper are illustrative examples of what the framework can host, not an exhaustive list. The architecture supports any constitutional parameterization that preserves the Model's load-bearing properties: equal per-citizen issuance at point of creation, formula-bounded magnitudes, separated banking, locked Stable Floor accounts for K1 and K2, and constitutional governance over rule changes. The consumer-dividend share κ_d (Section 3.4) is one such parameter — a dial, available to any Mode, that splits the real-growth budget between locked floors and a spendable citizen dividend. Within those constraints, a society's menu of valid Modes is whatever the citizens choose to ratify.

Other coherent Modes a society could ratify include, but are not limited to:

- A higher-inflation target Mode (e.g., 4 percent rather than 2 percent), appropriate during periods of unusually high real debt overhang.
- A variable-target Mode tied to unemployment or output gap, effectively a constitutional NGDP-targeting regime.
- An asymmetric KI Mode that activates only when inflation falls below target — a one-sided deflation stabilizer.
- A pure citizenship Mode in which K2 is inactive and only K1 operates, producing more aggressive deflation than Mode A.
- A higher-K1 endowment Mode in which the citizenship endowment is calibrated as a more meaningful share of lifetime earnings, with K2 reduced proportionally.
- A variant emphasizing even tighter price-level control than Mode B's steady state, with K2 dynamically matched to nominal GDP growth on a shorter adjustment cadence. This variant accepts tighter calibration sensitivity in exchange for a genuinely neutral price level, and is most appropriate for economies with the deepest long-duration nominal contract infrastructure.
- A KI-as-emergency-tool Mode operating within Mode A or Mode B that keeps KI permanently inactive in normal conditions but makes it available as a bounded, automatically-reversing emergency channel — similar to the existing emergency toolkit but operating through the direct-deposit transmission rather than through capital markets. This gives deflation-targeting and stability-targeting Modes a fast-acting circulating-pool injection tool without permanently activating KI, preserving the deflationary or stable price-level character of the base Mode in normal times.
- Adaptive multi-governor configurations — such as Mode Λ described in Section 9 — that combine demographic, productivity, and conditional KI governors into a single responsive system.
- A consumer-dividend Mode (positive κ_d), in which a constitutional share of the real-growth budget is paid to citizens as a standing, universal, immediately spendable dividend instead of being routed entirely into Stable Floors. κ_d is not a knob slid onto an existing Mode: a different consumer share is a different Mode, defined by its own calibration vector (the K1/K2/K3/KI settings together with κ_d) and ratified whole, because the channels are coupled on the growth budget and the price locus and must be vetted together. Because K3 draws from the same growth budget as K2 rather than adding to it, a κ_d held within the circulating-pool ceiling defines a Mode that preserves a base Mode's price regime — a Mode B with $\kappa_d > 0$, for example, keeps exact price stability while paying a current dividend — and a κ_d set beyond that ceiling defines a higher-inflation Mode. The parameter thus both spawns new Modes and supplies the governed path between them; it is the present-tense complement to the floor-building channels, trading future capital for current income at a rate the citizens ratify.
- A Mode 0, or zero-issuance Mode, in which every creation channel is inactive (K1 = K2 = KI = 0) and no new money is created at all — with real output still growing against a fixed supply, the price level falls at roughly the real-growth rate, recovering the fixed-supply, hard-money corner as one admissible setting of the same engine rather than a separate system. It is the limiting case that answers the reflexive

printing-press objection directly: zero is one of the framework's settings, not a press stuck on "more." Even here the citizen ends ahead of the status quo, because it is today's dollar that quietly loses value.

The constitutional ratification process for any new Mode is identical: 67 percent supermajority, 90-day deliberation, simulation-backed calibration, and FDCA verification that the proposed parameters preserve the Model's load-bearing properties. The paper presents five illustrative Modes: four base Modes — Modes A, B, and C spanning the inflation-deflation spectrum (Sections 4–6), and Mode D, which holds Mode B's price-stable regime but distributes the entire budget as a standing dividend with no locked floor — and one adaptive multi-governor configuration (Section 9). The full menu of options is not limited to five; the framework's claim is that any coherent monetary regime within the Model's bounds can be ratified by the same constitutional process.

8. Mode Ω — The Price-Stability Configuration

The base Modes issue on fixed recipes and let the price level fall where it may. Mode B, in particular, issues at the full growth rate on a fixed 60/40 split; the resulting inflation is an outcome, not a target. That outcome is near zero only for an economy whose transaction-active money share sits near one-half (Section 5). Mode Ω removes the fixed recipe and targets the outcome directly: it is the governed configuration that holds the derived price level flat in any monetary structure, by solving the split rather than fixing it. Like Mode Λ it is adaptive and formula-driven — no committee, no discretion — but where Λ responds to demographic and productivity conditions, Ω responds to the single objective of price stability.

8.1 The Two-Sided Stabilizing Rule

Mode Ω begins from the same floor-building commitment as Mode B — the 60/40 split of the real-growth-matched budget — and then drives the derived inflation to zero from whichever side it would otherwise miss. Where the fixed split would inflate (a thin transaction circuit, low μ), Ω trims the standing dividend κ_d , and the freed budget flows into the locked floor. Where the fixed split would mildly deflate (a thick transaction circuit, high μ), Ω adds the KI inflation-gap channel, injecting the shortfall as spendable dividend. The two levers are one-directional and complementary: the dividend trim can only reduce the transactional injection, the KI channel can only add to it, and exactly one is active in any given economy. The result is a price level held at zero across every calibrated monetary structure, rather than the mild drift the fixed split produces away from its balance point.

8.2 The Stabilizing Split Is Structural, Not Growth-Dependent

The split that achieves price stability has a closed form. Setting the derived inflation to zero on the inflationary side and neglecting the small fixed-floor term, the stabilizing

dividend share reduces to $\kappa_d^* \approx (\mu - 0.20) / 0.80$, where μ is the transaction-active share of broad money. The growth rate cancels: it enters both the money injected and the deflationary pull that injection must offset, and those scale together. Price stability is therefore a structural property of an economy's monetary plumbing — how much money sits in transaction accounts rather than time deposits — and is essentially independent of how fast the economy grows. For a transaction share near one-half the stabilizing split is near 40 percent (Mode B's fixed value, which is why Mode B is near-stable for such economies); for a time-deposit-heavy economy near 0.28 it falls to roughly 10 percent. What growth governs is not the split but the size of the budget: the same stabilizing share funds a far larger dividend and floor in dollar terms in a fast-growing economy than in a slow one.

8.3 Domain of Validity and the Contraction Hand-Off

Mode Ω holds the price level flat for all positive real growth. At zero or negative real growth its lever is exhausted, and this boundary is exact rather than approximate. When real growth falls to the point where the growth-funded budget collapses to zero, there is no dividend left to trim; and a contracting economy is inflationary, not deflationary, by the standard identity — a roughly fixed transactional stock against a shrinking real output raises the price level — so the KI channel, which only adds money, would worsen the gap rather than close it. Ω therefore does everything its split-governor allows (trimming the dividend to its floor) and then hands the residual to the emergency toolkit. This is not a failure of the mode but a correct statement of its domain: the split-governor governs issuance under growth; removing a contraction-driven inflation requires active money retirement, which is the province of the Emergency K2 Attenuation rung (Tool 14c) and, beyond it, the Surge Brake (Tool 14a). The point at which Ω 's dividend trim reaches zero is precisely the point at which Tool 14c arms.

8.4 Relation to Mode B, and a Rejected Alternative

Across the calibrated economies Mode Ω weakly dominates the fixed Mode B: for every economy the locked floor is at least as large and the total lifetime value at least as high, and for economies away from the balance point (thin- or thick-circuit) it is strictly better, holding prices flat where Mode B drifts. On the inflationary side the improvement is unambiguous — trimming the dividend both stabilizes prices and enlarges the locked floor. On the deflationary side, Ω relocates a small share of value from the locked, tax-advantaged floor into the liquid, spendable dividend via the KI channel; total lifetime value is preserved or higher, because the dividend stream compounds over the working life more than the marginal floor contribution it replaces. A wealth-first alternative was considered and rejected: an Ω that fought only inflation and tolerated the deflationary side would hold a marginally larger locked floor in thick-circuit economies, but it would forgo price stability there and would not hold the name the mode claims. Because persistent mild deflation carries a real cost through downward nominal wage rigidity — a cost the framework does not otherwise dismiss — and because the value it relocates is not destroyed but returned to citizens as spendable dividend, the two-sided stabilizing

form is the one specified. A polity that genuinely prefers mild deflation as a feature already has Mode A for that purpose; Ω is the configuration for a polity that wants the price level to hold.

8.5 The Ownership Dial

Because Mode Ω stabilizes the price level through the split, it leaves a second quantity free: the share of the equity market that the locked floor accumulates on citizens' behalf. That share can be set to any target at or below its natural level without disturbing price stability, by trading locked index purchase (the K2 channel) for liquid dividend (the KI channel). The two objectives do not compete. Given an ownership target, the required index purchase fixes K2; holding the price level flat then fixes KI, because lowering K2 opens a deflationary gap that KI fills exactly. The system is triangular and always closes when KI is non-negative, which holds for every target at or below the natural level. Ownership thus becomes a policy dial rather than a fixed consequence of the design, and the value a citizen would have held as equity instead arrives as spendable dividend: a change in the form of the citizen's benefit, not its magnitude.

The cost of holding a target share is smaller than intuition suggests. To maintain a steady ownership share, the citizen circuit need only absorb that share of net new equity issuance each period, not that share of total market growth, because the citizen stake and the market appreciate together and appreciation cancels in the ratio. Net issuance is a small flow in most markets, so even a large ownership share imposes only light purchase, and the target tunes it downward from there. The full derivation, the closed form, a worked calibration, and the connection to the structural buyer's fixed point (where absorption equals net issuance) are given in the Structural Buyer paper (Neo-Solon 2026h, Section 6.5 and Appendix C). The practical significance is that a concern often raised against a universal citizen stake, that it would accumulate too large a share of the market, is answered structurally: the share is bounded by rule, at zero cost to price stability.

9. Mode Λ — An Adaptive Configuration

Mode Λ is an illustrative configuration that demonstrates how the framework's modular governor architecture can be combined into a single adaptive system. Like Modes A, B, C, and D, it is a constitutional parameterization of the Model — not a recommendation, but a coherent and fully specified option that a society could ratify through the standard Tier 2 supermajority process. It is presented here alongside the four base Modes because the paper's claim — that the framework can host any coherent monetary regime within the Model's bounds — is more credible with a concrete adaptive example than as an abstract assertion.

Before describing Mode Λ 's mechanics, it is useful to locate it in the comparative landscape established in Section 7. Mode Λ does not occupy a fixed position on the inflation-deflation spectrum the way the base Modes do. Instead, it occupies an adaptive band whose position responds to observable conditions. The following table shows

where Mode Λ sits relative to Modes A, B, and C across its key parameters — in normal conditions and under the demographic stress scenario for which its governors are specifically designed.

Metric	Mode A	Mode B	Mode C	Mode Λ (normal)	Mode Λ (stress)
Inflation target	-1.86%	0% drift	+2%	\approx -0.8% (mild deflation)	Up to -0.05% (soft floor: -1.2%)
Stable Floor at 65 (real)	\sim \$233K	\sim \$413K	\sim \$230K	\sim \$403K	\sim \$343K–\$638K
K1 multiplier	1.0 \times (fixed)	1.0 \times (fixed)	1.0 \times (fixed)	1.0 \times (inactive)	Up to 2.0 \times (demographic)
K2 booster	None	None	None	None (60% base)	Up to +40% (demo + productivity)
KI stream	Inactive	Inactive	Permanent (+2% target)	Inactive	Conditional (deflation floor)
Inflation governed by	Fixed formula	Fixed formula	Fixed formula	Fixed formula	Adaptive formula
Ratification process	Tier 2 supermajority	Tier 2 supermajority	Tier 2 supermajority	Tier 2 supermajority	Tier 2 supermajority

Table 7. Mode Λ vs. Modes A, B, and C. 'Normal' reflects base-case conditions (pop growth +0.3%, central general-equilibrium return) where governors are largely inactive. 'Stress' reflects negative population growth (-0.5%) with the demographic governor active. Stable Floor figures from Table 8 (Section 9.6). All Modes share the same ratification process. Mode Λ 's variable Stable Floor range reflects scenario dependence rather than a single deterministic baseline.

The key comparative insight is this: in normal demographic conditions, Mode Λ behaves conservatively — a mild deflationary bias from its 60-percent capture rate, no KI dividend, conservative capital stake accumulation. Its differentiation from Mode B emerges precisely when demographic or productivity stress arrives. A society that expects stable population growth and wants simplicity is better served by Mode A or Mode B. A society that wants automatic protection against demographic stress — without relying on emergency tools or constitutional amendments to respond — has a reason to consider Mode Λ . That is the tradeoff, stated directly.

Mode Λ synthesizes elements from the four base Modes and from demographic-adaptive and productivity-responsive configurations. Its distinguishing feature is that it combines multiple formula-driven governors into a single system, rather than running any one governor in isolation. Every multiplier, threshold, and activation trigger in Mode Λ is formula-derived from publicly published data and constitutionally specified in advance — adaptive means the issuance formula responds automatically to observable conditions, not that a committee exercises judgment about how to respond. The distinction from discretion is total: no human decision is required at activation, and no human decision can override the formula.

Mode Λ at a Glance

- What it is: A multi-governor configuration combining adaptive K1 multipliers, K2 boosters, and a conditional KI stream in a single constitutionally specified formula.

- What it does: Automatically increases citizen issuance when population growth turns negative, when productivity surges above trend, or when deflation triggers the KI safety valve — then reverts to baseline once conditions normalize.
- Inflation rate: no fixed target. Mildly deflationary in normal conditions ($\approx -0.80\%$, the 60-percent-capture base); governors ease the deflation toward 0% under stress, reaching $\sim -0.08\%$ at the productivity-boom peak. KI creates a soft floor preventing deflation below -1.2% . (See Figures 5a–5b.)
- What it is not: Discretionary. Every activation threshold, multiplier range, and reversion rate is formula-specified and publicly auditable. No committee judgment is involved.
- Why it is here: As a worked example showing the Mode menu is genuinely extensible — not as a recommendation. Societies may prefer a simpler Mode.
- Key limits: K1 capped at $2.0\times$ base; combined issuance capped at 3.5% of M2; KI carries a 36-month sunset with mandatory reaffirmation vote. (Full guardrails: end of this section.)

9.1 Core Goal

Maximize long-term citizen prosperity by automatically balancing: price stability with mild deflationary bias; demographic resilience (especially aging and negative population growth); productivity surges such as those produced by AI and automation; and reasonable current-life support without gutting retirement compounding.

9.2 K1 — Citizenship Endowment (Adaptive)

Base calibration: 2.5 percent of GDP per capita at the moment of citizenship, identical to the base Modes.

Adaptive multipliers (stackable, capped at $2.0\times$ total):

- Demographic multiplier: $+0.3\times$ to $+0.8\times$ based on old-age dependency ratio and total fertility rate trends. Activates when the dependency ratio rises above its 10-year baseline or when TFR falls below replacement.
- Population growth multiplier: additional boost when net population growth falls below $+0.2$ percent annually, with a stronger boost in negative population growth territory.
- Youth bonus: $+20$ percent for births to citizens (front-loads human capital investment by weighting the citizenship endowment toward new births relative to naturalizations).

The $2.0\times$ total multiplier cap is derived from the constraint that total K1 issuance under maximum multipliers cannot exceed 5.0 percent of GDP per capita — twice the base calibration — which keeps K1 issuance within the range that the Adaptive Smoothing Rule can absorb without producing instability in the Stable Floor accumulation trajectory.

9.3 K2 — Growth Dividend (Adaptive)

Base calibration: approximately 60 percent of real-growth-matched, placing Mode Λ between Mode A's 19.5 percent (of M2 growth) capture rate and a full real-growth

match. This produces mild deflationary bias under normal conditions while preserving room for upward adjustment through the adaptive boosters.

Adaptive boosters:

- Productivity Governor: +10 to +30 percent of base K2 when the Composite Productivity Index exceeds its 5-year trend by more than one standard deviation. This channels productivity windfalls directly into Stable Floor accounts.
- Demographic Governor: +15 to +40 percent of base K2 during periods of rising dependency ratio or negative population growth. This partially compensates for the reduced per-cohort compounding window that smaller future generations face.

9.4 KI — Supplemental Citizen Stream (Conditional)

KI in Mode Λ is conditional rather than permanently active. It activates only when either: (a) deflation exceeds 1.2 percent annually for two consecutive quarters, or (b) demographic stress (rising dependency ratio) coincides with below-trend wage growth for three consecutive quarters. When activated, KI issues at 0 to 0.6 percent of M2 annually, split 60 percent to Stable Floor accounts and 40 percent as an immediately spendable citizen dividend. KI automatically deactivates when the triggering conditions resolve.

9.5 Additional Built-in Features

- Extended Adaptive Smoothing Rule — the 5-year rolling average window extends to a 7-year window during periods of elevated volatility, reducing sensitivity to single-year data anomalies.
- Automatic reversion to conservative baseline after temporary shocks — any governor above baseline reverts at 25 percent per year after the triggering condition resolves, preventing cliff effects.
- Full transparency: all multipliers published monthly via public dashboard using independently verified data from the same five agencies that feed the Composite Productivity Index.

Mode Ω — How the Governors Respond to Conditions Over Time
 (Illustrative 65-year environment: aging population, mid-century productivity boom, late-period deflation)

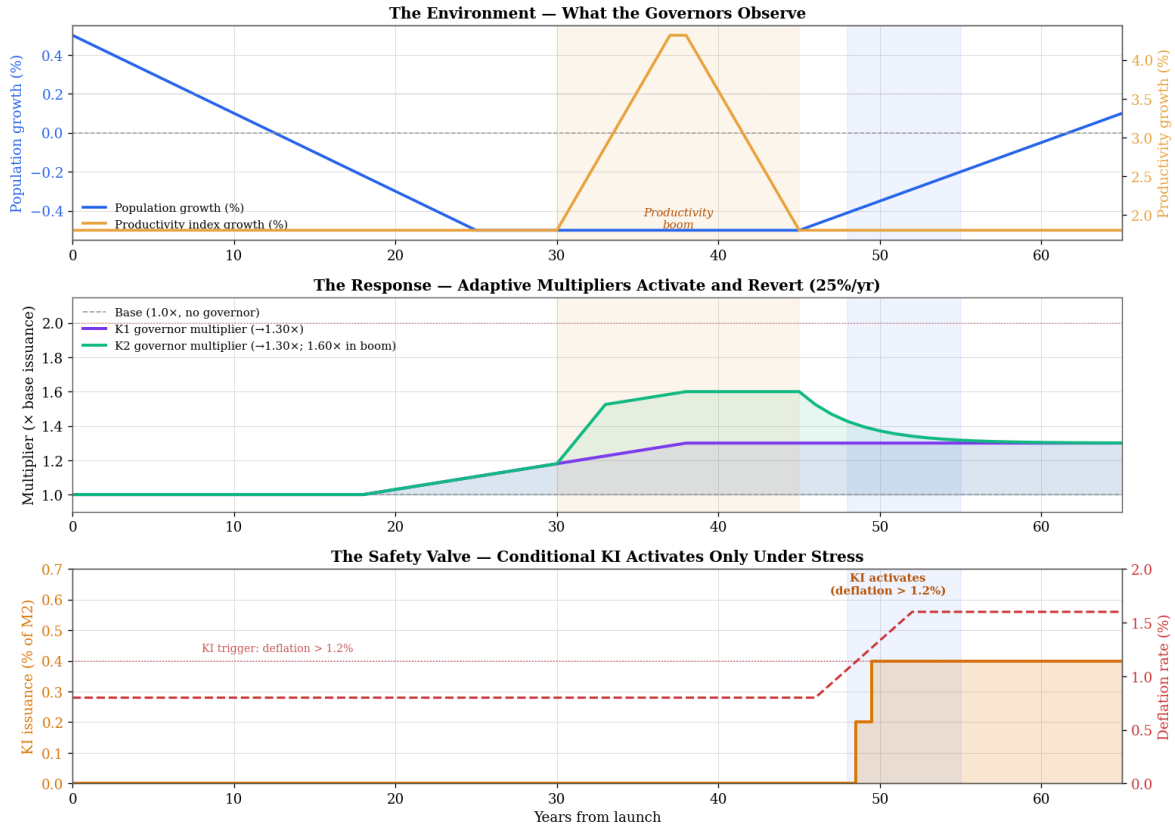


Figure 4. How Mode Λ 's governors respond to conditions over an illustrative 65-year environment. Top: the inputs the governors observe (declining population growth, a mid-century productivity boom). Middle: the K1 and K2 multipliers activating in response and reverting toward baseline at 25% per year. Bottom: the conditional KI stream activating only under stress — partially during demographic stress with low wage growth, fully when deflation exceeds 1.2%. Every governor is formula-driven from published data; no committee judgment is involved.

9.6 Simulation Results — 2025-Born Citizen at Age 65 (in 2025 dollars)

Scenario	Final Stable Floor	Annual Income (5% cap)	Notes
Normal pop growth (+0.3%), GE realizable return	~\$403,000	~\$20,200	Base — multipliers largely inactive
Negative pop growth (-0.5%), central GE return	~\$480,000	~\$24,000	+19%; governor lifts capture, return attenuates
Negative pop + high α/δ band	~\$638,000	~\$31,900	Excellent outcome
Negative pop + low α/δ band	~\$343,000	~\$17,200	Still respectable
Productivity boom + negative pop	~\$541,000	~\$27,000	Very strong

Table 9. Mode Λ illustrative simulation results for a citizen born at launch, age 65 in 2090, in 2025 dollars. These are deterministic baselines on the general-equilibrium realizable return (Neo-Solon, 2026e §6.7); the companion empirical paper (Neo-Solon, 2026b) provides stochastic characterization. Mode Λ 's base-case Stable Floor sits alongside Mode B's and well above Modes A and C, reflecting its 60-percent base K2 capture; the adaptive governors provide additional accumulation precisely when and where it is most needed.

9.7 What Mode Λ Demonstrates

Mode Λ demonstrates three properties of the framework that the four base Modes do not individually illustrate:

- The governor architecture is composable. K1 multipliers, K2 boosters, and a conditional KI can be combined without violating any of the Model's seven load-bearing properties. Each governor is still formula-derived, citizen-anchored, and publicly auditable.
- Demographic resilience can be built into the Mode rather than addressed through emergency tools. The base Modes respond to demographic stress through the emergency toolkit (Section 12); Mode Λ demonstrates that demographic adaptation can instead be constitutionally specified in the Mode's calibration formula.
- The Mode menu is genuinely open. A society that wants neither Mode A's pure deflation nor Mode C's permanent dividend nor Mode B's conservative stability can ratify a configuration like Mode Λ — or design something entirely different — through the same Tier 2 supermajority process.

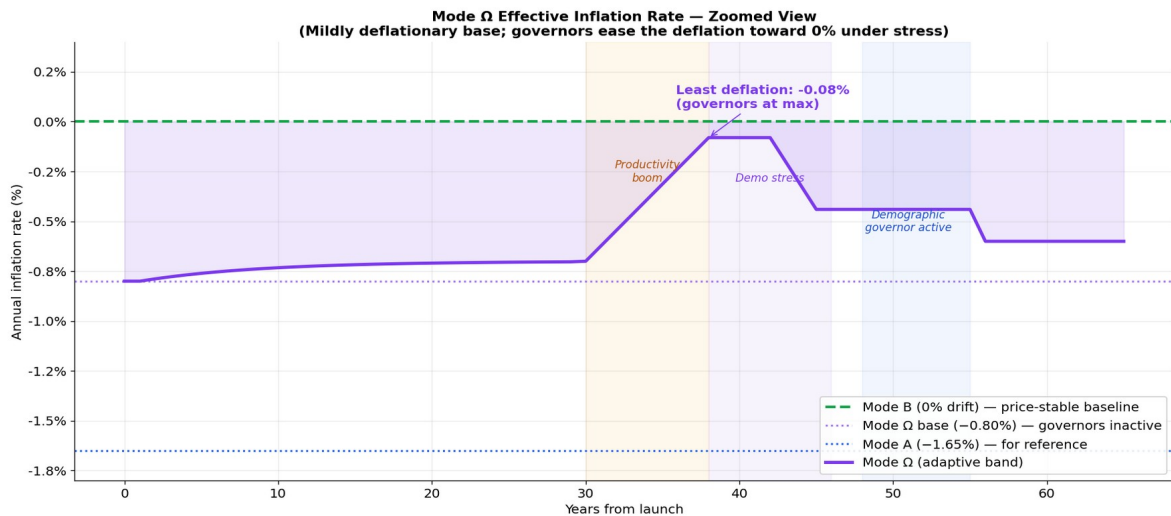


Figure 5a. Mode Λ effective inflation rate — zoomed view (Y-axis cropped to show the adaptive band). Mode Λ runs at its mildly deflationary 60-percent-capture base of approximately -0.80% in normal conditions; when the demographic and productivity governors engage they add issuance and ease the deflation toward 0% , reaching approximately -0.08% at the productivity-boom peak — never crossing into positive inflation — before reverting toward the -0.80% base. Governors push issuance higher through the capital-markets channel (minimal direct CPI effect). KI activates as a soft floor only if deflation would otherwise breach -1.2% . Mode C ($+2\%$) and the current system ($+3\%$) are far above this range and not shown.




Figure 5b. Mode Λ inflation gap below the Mode B baseline during governor-active periods. Mode Λ sits approximately 0.80% below Mode B at its base; the governors narrow the gap rather than widen it, reaching approximately -0.08% (near price stability) when demographic stress and the productivity boom coincide. Once conditions normalize the gap reverts toward the -0.80% base. The result is an adaptive band rather than a fixed inflation target: Mode Λ is mildly deflationary in calm periods and eases that deflation toward 0% when stress conditions warrant additional issuance.

Mode Λ is not a recommendation. A simpler society might find Mode A or Mode B more governable precisely because their formulas are easier to verify and harder to misinterpret. Mode Λ is offered as evidence that the framework’s claim — that any coherent monetary regime within the Model’s bounds can be constitutionally ratified — is not merely theoretical.

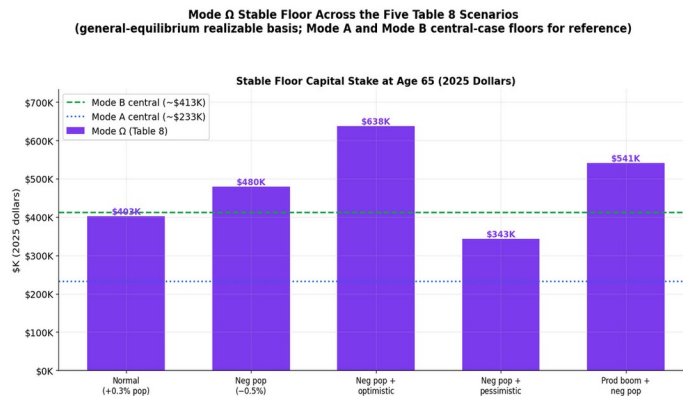
9.8 Constitutional Guardrails

Any society ratifying Mode Λ or a comparable adaptive configuration should specify, alongside the governor formulas, a set of constitutional guardrails that address the unique risks adaptive Modes introduce. The following are recommended as Tier 2 provisions in any Mode Λ ratification.

- Governor floor: the base K2 calibration rate (60 percent of real-growth-matched) cannot be reduced below 40 percent by governor activation alone. The Demographic Governor provides boosts above baseline; it cannot substitute for legitimate K2 issuance by drawing down the base rate.
- Multiplier review trigger: if the K1 demographic multiplier remains above $1.5\times$ for more than 15 consecutive years, a mandatory FDCA review is triggered to determine whether the multiplier is responding to genuine temporary demographic stress or whether the underlying data inputs have drifted in a way that makes the governor permanently active. This distinguishes a stressed environment (correct response) from a calibration error (incorrect response).
- KI sunset: each activation of the conditional KI stream carries an automatic 36-month sunset. If the triggering conditions persist beyond 36 months, KI continuation requires a supermajority reaffirmation vote — not a full Mode ratification, but an explicit citizen decision to continue the activation. This prevents KI from becoming

permanently active through sustained demographic stress, which would effectively transform Mode Λ into Mode C without a formal Mode change vote.

- Governor interaction cap: when both the K1 demographic multiplier and K2 demographic booster are simultaneously active above their minimum thresholds, combined total issuance (K1 + K2 + any active KI) cannot exceed 3.5 percent of M2 annually without a Tier 2 amendment. This prevents a scenario where all governors activate simultaneously and produce more circulating-pool expansion than Mode C would in steady state.
- Amendment process for governor recalibration: changing the multiplier ranges, governor thresholds, or reversion rates requires a Tier 2 amendment by the same 67 percent supermajority as any other Mode change. The FDCA may propose recalibration after the mandatory review trigger, but implementation requires citizen ratification.



Mode Ω runs a 60-percent base capture with adaptive demographic and productivity governors. At the central return it sits just below fixed-capture Mode B; at the low end of the return band the demographic governor outweighs Mode B (Mode Ω leads Mode B's roughly \$277K there), and it leads most clearly when the productivity governor also engages.

Figure 6. Mode Λ Stable Floor at age 65 across the five demographic and equity-return scenarios of Table 8 (Section 9.6), on the general-equilibrium realizable basis, with the Mode A and Mode B central-case floors (\approx \$233K and \approx \$413K) shown for reference. Because Mode Λ runs a 60-percent base capture with adaptive demographic and productivity governors, its floor sits between Mode A and Mode B at the central return; under demographic stress at the central return it approaches Mode B (\approx \$480K), at the low end of the return band it exceeds Mode B (\approx \$343K versus \approx \$277K, where the demographic governor outweighs Mode B's fixed capture), and it leads most clearly when the productivity governor also engages (\approx \$541K). The scenario floors are the published Table 8 values; the macroeconomic model (Neo-Solon, 2026e, Section 6.7) supplies the realizable return per scenario, including the capture-return feedback through which lifting capture to defend the floor attenuates the return it earns.

10. Mode T — The Transition Configuration

The four base Modes and Mode Λ describe steady-state operation. They do not by themselves resolve a distinct one-time problem: how a society carrying a large pre-existing sovereign debt migrates from the current system to the Citizens Standard without either defaulting, inflating the debt away through a consumer-price shock, or

diverting citizen wealth accumulation to debt service. Mode T is a purpose-built configuration for exactly this migration. It is transition-only and self-extinguishing; it is presented here as a named example of the architecture's flexibility, in the same spirit as Mode Λ , rather than as a permanent operating regime.

The configuration. Mode T runs three channels simultaneously. K1 operates at its standard calibration, depositing the citizenship endowment into each new citizen's Stable Floor. K1 and K2 together operate at the full real-growth-matched rate, with K2 the residual after K1, so that new money creation matches real output growth. For an economy at the one-half transaction-ratio balance point (the United States is close, and is the worked case throughout this paper) the fixed split holds the price level flat with the citizen-dividend channel KI inactive. For an economy away from that point the same fixed split drifts, so the transition runs the Mode Ω solve (Section 8): KI is active and the split is set so that derived inflation is zero throughout the transition, not merely at its end. In both cases the price level is held stable while the transition proceeds. In its place, the transition-only channel KT issues money calibrated to a price-level path, directed not to citizens but to the Legacy Debt Trust for retirement of pre-existing public debt. The mechanics of the Trust and the KT channel are specified in the transition paper (Neo-Solon, 2026c) and the statutory implementation paper (Neo-Solon, 2026d). The macroeconomic model (Neo-Solon, 2026e, Section 4.7) further specifies an optional conditional damper, KI_T, that absorbs the residual credit contraction the transition lending facility does not fully cover, as additive net-new issuance that leaves K1 and K2 untouched and self-extinguishes once banking separation completes.

Why KT is a separate channel. KT differs from KI in every operationally relevant respect: its purpose is debt retirement rather than citizen distribution; its recipient is the bondholder receiving redemption rather than the citizen receiving a dividend; its economic signature is an asset swap (a bondholder exchanges a bond for cash and predominantly reinvests) rather than new consumer demand; and its lifecycle is bounded rather than permanent. Because the redeemed funds are overwhelmingly held by structural reinvestors — foreign central banks, pension funds, and the central bank itself — the consumer-price impact of KT is minimal, and the channel is calibrated to a price-level path so that issuance automatically throttles if consumer inflation were to rise. KT therefore retires debt without the consumer-price shock that would accompany distributing the same magnitude of new money to households.

Consumer-price neutrality, asset-price honesty. Mode T holds consumer prices stable through two facts: KT is a price-neutral asset swap that redeems bonds and never enters the transaction circuit, and the citizen-facing channels are configured for zero derived inflation, by the fixed split where the economy sits near the one-half transaction ratio and by the Mode Ω solve (Section 8) otherwise. Price stability during the transition therefore does not depend on the economy's transaction ratio. It is not, however, neutral with respect to asset prices: bondholders receiving redemption reinvest, and as sovereign instruments become scarcer a portion of that capital rotates into equities and real estate, exerting mild upward pressure on asset valuations. This effect is stated plainly rather than minimized. Its principal incidence is benign for the framework's

purpose: because citizen Stable Floors hold broad-market equity, the appreciation accrues substantially to citizens rather than to a narrow holder class. The quantitative treatment, including its modest intergenerational return-compression effect, is given in the transition paper.

10.1 Mode T-stable — The Automatic Steady State

The self-extinguishing design. KT is bounded by construction. As public debt enters its operational band (specified in the transition paper at approximately 30 to 60 percent of GDP), the KT channel throttles down automatically, retiring only what is needed to hold the band and routing the remaining growth-matched seigniorage to citizen Stable Floors. No vote is required; it is a parameter of the channel itself. The configuration that remains — K1 and full-rate K2, with sixty percent funding Stable Floors and forty percent paid as the standing K3 dividend (the 60/40 split), KI inactive, and KT dormant within the band but available as a symmetric open-market instrument — is Mode T-stable.

Continuity guarantee. The transition from Mode T to Mode T-stable is, from the citizen's perspective, a non-event. K1 and K2 flow identically before and after the threshold is crossed. Because KT never deposited into citizen accounts — it redeemed bonds — its deactivation produces no discontinuity whatsoever in Stable Floor accumulation. There is at no point an undefined monetary state: the system is always in a fully specified configuration, and the default configuration after debt stabilization is itself permanently sustainable.

Mode T-stable is a sound permanent regime, and it is Mode Ω . When KT sunsets, the citizen-facing configuration that remains is Mode Ω , the price-stability steady state of Section 8: money creation matched to real output growth, the split solved so derived inflation is zero, and KI available to correct drift in either direction. For an economy at the balance point this coincides with the fixed Mode B split; for any other economy it is the Ω solve that Mode B is not. Mode T-stable therefore delivers stable prices for any transaction ratio, indefinitely, without further intervention and without requiring that any debt exist. It is not a temporary bridge awaiting a decision; it is a complete monetary system that can operate permanently.

The optional steady-state choice. Once in Mode T-stable, a society may — but need not — choose a different permanent configuration through the standard Tier 2 supermajority process, now informed by decades of observed performance rather than projection. It may remain in Mode T-stable (the default, requiring no action); activate KI as a citizen dividend in the manner of Mode C, the channel capacity having been freed by the retirement of the debt; reduce K2 below the growth-matched rate to permit the mild deflationary accumulation of Mode A; or hand the parameters to the adaptive governors of Mode Λ . Each destination is itself stable. The transition thus has a guaranteed soft landing regardless of the political outcome of any subsequent mode vote.

11. Banking Architecture

The Citizens Standard restructures banking around a single principle: the payment system is a public utility and must never fail; the credit system is a market activity and may fail in ordinary ways without bringing the payment system down with it. The banking architecture is identical across all Modes.

11.1 The Payment-Credit Separation

Two account types replace the current deposit structure. Transaction accounts are full-reserve. Funds in transaction accounts are held one-for-one with FDCA reserves and are constitutionally protected from any creditor claim, any bank insolvency, and any government seizure. They are functionally identical to physical cash, but in digital form.

Term deposits are at-risk investment claims, not money. Citizens may choose to deposit funds for fixed terms at competitive interest rates. Term deposits explicitly fund bank lending and explicitly carry credit risk. The depositor knows this, accepts it, and is compensated through interest. There is no FDIC-style guarantee on term deposits. This is the Chicago Plan architecture, validated empirically by the IMF in 2012, with the modification that term deposits provide an explicit voluntary mechanism for citizens to channel savings into productive lending.

11.2 How Credit Is Created

Banks fund loans entirely from term deposit liabilities and equity capital. They cannot create new money through lending. The protocol sets the maximum leverage at 4:1 under normal conditions. The leverage ratio is countercyclical — when private credit growth exceeds nominal GDP growth by more than three percentage points for two consecutive quarters, the maximum leverage automatically tightens to 3:1; when private credit growth is more than three percentage points below nominal GDP growth, leverage loosens to 5:1.

11.3 The Transition Lending Facility

The transition from current banking architecture creates a substantial credit contraction risk. The framework solves this through a Transition Lending Facility (TLF) operated for thirty years from implementation. The TLF is funded from a fraction of citizen seigniorage retained from K2 issuance during the transition. Funds route to community banks, agricultural lenders, and small business lenders at zero spread. The TLF dissolves on schedule, with lending volume declining linearly to zero over the thirty-year transition.

11.4 Mode-Specific Banking Considerations

Banking architecture is identical across Modes, but the inflation regime affects credit pricing. Mode A's deflation target requires universal price-indexed debt. Mode B's stability target allows conventional nominal contracts to function natively. Mode C's 2 percent inflation target produces conventional nominal contract behavior similar to the

current system, with the substantial difference that the 2 percent inflation rate is constitutionally specified and formula-stabilized. Mode Λ's banking considerations follow whichever inflation outcome its governors produce in any given period.

12. Emergency Tools and Crisis Response

A monetary framework that cannot respond to crises will be abandoned during the first serious crisis it faces. The Citizens Standard provides a suite of sixteen bounded, rules-based, automatically reversing tools that preserve the system's integrity while providing genuine crisis response capacity. The toolkit is constitutionally available across all Modes. Every tool satisfies the same three requirements: rules-based (triggered by objective criteria, not judgment), bounded (hard constitutional ceiling), and automatically reversing (expires without active renewal). The tools are organized into four failure-mode categories so that responses match the actual shape of the crisis.

12.1 Failure Mode 1 — Demand Collapse

Demand collapse is the failure mode of insufficient spending: velocity falls, prices fall or stagnate, real GDP contracts, unemployment rises. The right response is to support spending power directly.

Tool 1 — Emergency K1 Provision. K1 ceiling rises temporarily to 2.0 percent of M2. Triggered when deflation exceeds 1 percent for two consecutive quarters after K2 has already ceased. Ceiling: 2.0 percent of current M2 (\$447B/yr at launch M2 of \$22.4T). 24-month automatic sunset.

Tool 2 — Constitutional Rainy-Day Fund. Surplus reserve deployed to enhanced unemployment insurance, equal per-capita direct payments, and healthcare cost coverage only. Triggered by unemployment over 8 percent for 60 consecutive days, or real GDP contraction over 3 percent annualized for two quarters. Ceiling: 3.5 to 4.5 percent of current M2 (\$795B–\$1.02T at launch) one-time. Operates on existing reserves; no new money created.

Tool 4 — Emergency Fiscal Reallocation. Up to 30 percent of pre-specified non-essential discretionary categories redirected to unemployment insurance, equal per-capita payments, and healthcare. Ceiling: 1.5 to 2.2 percent of current M2 (\$340B–\$500B at launch). Operates on existing fiscal flows; no new money created.

Tool 7 — Immediate Velocity Dividend. Equal per-capita immediate cash payment to all verified citizens. Triggered by productivity growth over 3 percent for two consecutive years. Ceiling: 0.10 percent of M2 (\$22.4B/yr).

Tool 8 — Deflationary Floor Dividend. Equal per-capita new money to all citizens, immediately spendable. Triggered when net population growth is negative for two quarters AND real GDP under 0 percent for two quarters AND velocity index more than 10 percent below its 36-month average. Ceiling: 0.05 percent of M2.

Tool 9 — Enhanced Automatic Stabilizers. Unemployment insurance, food assistance, Medicaid expand 25 percent above statutory baseline automatically when unemployment exceeds 6 percent for 30 consecutive days.

Tool 10 — K2 Emergency Boost. K2 capture rate temporarily increases (within constitutional bounds) after recovery begins. All goes to locked Stable Floor accounts — zero circulating inflation.

Tool 11 — Stable Floor Bridge Loan. Citizens 60+ may borrow up to 10 percent of account value at 0 percent interest against their own equity (not a withdrawal). 5-year repayment from account or estate.

Tool 15 — M2 Contraction Floor. Addresses the Fisher debt-deflation cascade specifically — the scenario in which credit losses force lending contraction, which shrinks term deposits, which contracts M2, which depresses asset prices and triggers further defaults. Triggered when M2 contracts more than 5 percent over any rolling 12-month window. The FDCA creates money up to a ceiling of 3.0 percent of current M2 per year (\$671B at launch) and distributes it as an equal, immediately spendable per-citizen direct deposit, placing purchasing power directly into circulation to arrest the contraction. Automatic 18-month sunset; continuation beyond the sunset requires affirmative FDCA certification that the trigger condition persists, which prevents the tool from quietly becoming permanent monetary expansion. The dynamic cascade model in the replication package shows Tool 15 reduces acute-cascade M2 contraction by approximately 3 to 4 percentage points; its sunset bounds its contribution in prolonged multi-year contractions by design. This is the demand-collapse tool of last resort: bounded, rules-triggered, and aimed at the one failure mode — a liability-side monetary collapse — that the other tools address only indirectly.

These three money-creating demand-collapse tools form a graduated ladder rather than overlapping authorities. **Tool 1 (Emergency K1 Provision)** is the mild first rung — a temporary lift of the standing K1 ceiling when deflation merely exceeds 1 percent. **Tool 8 (Deflationary Floor Dividend)** is the targeted middle rung — a small, immediately-spendable per-capita issuance for the specific conjunction of negative population growth, contracting output, and collapsing velocity. **Tool 15 (M2 Contraction Floor)** is the severe top rung — the last-resort response to a liability-side Fisher cascade, sized an order of magnitude larger (3 percent of M2) and triggered only by an actual 5-percent M2 contraction. Each rung has a strictly higher trigger and a strictly larger ceiling than the one below it, so the response scales with the severity of the collapse and no two rungs do the same job at the same time. Tool 16, the Demurrage Backstop, sits outside this ladder: it neither creates nor drains net money, but applies a small revenue-neutral carrying charge that restores velocity in a deflation-targeting Mode without abandoning the deflation target — the one failure mode the money-creating rungs cannot address, since issuing money to lift velocity would destroy the very price path the Mode is chosen to deliver.

Tool 16 — Demurrage Backstop. A Mode-specific velocity stabilizer for deflation-targeting configurations (Mode A and any custom deflationary Mode), developed in

Section 4.5. A small carrying charge is applied to transactional balances, offsetting the real return that mild deflation confers on idle cash, so that money continues to circulate without lifting the price path. The Stable Floors are exempt, as with the Surge Brake, and the charge applies only to the transactional circuit. Triggered when the velocity index falls more than 10 percent below its 36-month average while the Mode's deflation target is being met — the signal that benign productivity deflation is tipping into precautionary hoarding, distinct from the demand-collapse trigger of Tools 1, 8, and 15, which fire on contracting output. Ceiling: the charge is capped at 4 percent per year on transactional balances, the national-scale leakage bound above which holders flee into untaxed stores of value (Keynes's objection to Gesell); the effective rate is set near the realized deflation rate, on the order of 2 percent. Revenue-neutral by construction: proceeds are redistributed equally per capita, so no net money is created or destroyed and the tool adds nothing to the framework's money-creation capacity — only the relative incentive to hoard is removed. Automatic reversal as velocity returns toward its average. In a deep deflation-expectations spiral the leakage ceiling binds before the charge can fully offset the hoarding, so the backstop is a damper deployed alongside the money-creating tools, not a standalone clamp. The calibration against the money-demand and demurrage-experiment evidence is released as the demurrage modules in the replication package.

12.2 Failure Mode 2 — Inflation Surge

Inflation surge is the opposite failure mode. The framework's first response is structural and automatic. In all Modes: K2 ceases entirely when real growth is zero or negative. K1 reduces through the price-level adjustment formula. In Mode C specifically: K1 contracts automatically because actual_CPI exceeding target_CPI drives the formula toward zero. Mode Λ 's conditional K1 contracts by the same mechanism. This is the strongest automatic counter-inflationary mechanism in any Mode — Mode C and Mode Λ self-correct without any tool activation.

Tool 14 — Two-Tier Surcharge (Surge Brake / Anchor-Keeping Surcharge). Tool 14 withdraws circulating money by retiring it through a graduated surcharge, not by absorbing it into interest-bearing debt. The distinction is load-bearing: retired money accumulates no stock, so the mechanism carries no roll-over interest and no $r < g$ sustainability condition, and its cost is a bounded, non-compounding annual drag proportional to the inflation gap rather than a liability that can run away. Bond sterilization — issuing interest-bearing reserve certificates to absorb dollars — is the costlier alternative the framework deliberately does not use.

Tool 14a — Surge Brake. A surcharge that withdraws and retires circulating money when consumer inflation runs hot. Trigger is Mode-specific: in Mode A, CPI exceeds 5 percent for two consecutive months (approximately 6.9 percentage points above the -1.9 percent target); in Mode B, CPI exceeds 5 percent for two months (5 percentage points above target); in Mode C and Mode Λ , CPI exceeds target plus 3 percentage points for two months, preserving the same gap-from-target threshold across all Modes. A velocity trigger runs in parallel: the brake also activates when the velocity index exceeds its 36-month average by more than 15 percent for 30 consecutive days, whichever fires first —

so a confidence-driven flight from money is met at the speed it develops, not only after consumer prices confirm it. Ceiling: 3 percent of current M2 (\$671B at launch). Reversal is automatic and continuous as inflation returns toward target; because nothing accumulates, there is no stock to unwind.

Tool 14b — Anchor-Keeping Surcharge. A graduated, bounded surcharge that holds realized inflation inside a narrow anchor band, active only for a polity that has adopted the external-interopability commitment — the pledge to hold inflation at a common near-zero anchor for cross-border predictability, developed in the companion external-interopability paper. It engages earlier and at a lower threshold than 14a — when realized inflation exceeds the band ceiling for two consecutive quarters — and withdraws money sized to the gap above the band. Like 14a it only subtracts and only retires; the KI channel only adds, and the K1, K2, and K3 channels are untouched. Its bound is the same 3-percent-of-M2 per-period ceiling, so it adds no contraction authority beyond 14a. A domestic-only polity that prefers to let mild productivity-driven drift pass need not activate 14b — adopting the external standard is what switches it on. Honest scope limit: 14b neutralizes a monetary inflation gap; a persistent real imbalance (production short of absorption) would require the surcharge to run continuously, and its limit is then the real and political tolerance for that standing drag, not a divergent stock. At that boundary the anchor yields to real adjustment.

Tool 14c — Emergency K2 Attenuation. The inflation-side mirror of Tool 1. Where Tool 1 fights mild deflation by temporarily raising the K1 issuance ceiling, Tool 14c fights mild inflation by temporarily attenuating the K2 issuance rate — turning the money tap down rather than clawing money back. It is triggered when realized inflation exceeds the Mode's target by more than 1 percent for two consecutive quarters, the symmetric counterpart of Tool 1's deflation-exceeds-1-percent trigger, and it is universally available: unlike the Anchor-Keeping Surcharge (14b), it does not require the external-interopability commitment. The attenuation is graduated and bounded, sized to the overshoot, with a 24-month automatic sunset and automatic reversal as inflation returns to target; because it acts on the issuance flow and not the money stock, nothing accumulates and there is no stock to unwind. It sits below the Surge Brake (14a) in the graduated ladder: attenuation is tried first, and because it can only reduce issuance to zero and no further, its floor is self-defining — if issuance has been attenuated to zero and inflation persists (the contraction case, where a shrinking denominator lifts prices while the growth budget has already vanished), that exhaustion is precisely the condition that escalates to 14a's money retirement. Tool 14c closes the toolkit's remaining asymmetry — the deflation side carried a mild first rung (Tool 1) available to every polity, while the inflation side previously jumped from tolerance straight to the 5-percent crisis brake unless the external anchor had been adopted — and it provides the domestic landing point for a price-stability Mode whose split-governor lever is exhausted at or below zero real growth.

Read together, the inflation response is a single two-sided throttle on the transactional layer. On the issuing side, KI is the only channel that adds transactional money, and it contracts automatically toward zero as `actual_CPI` exceeds `target_CPI`. On the

withdrawing side, once issuance has reached zero, Tool 14 retires transactional money below that floor — KI dials the inflow down to zero, Tool 14 takes it negative. Both act only on circulating money, the price-relevant aggregate, while the locked Stable Floors (K1, K2) are never touched — which is why the throttle can be this aggressive without threatening the run-proof lock. Tool ceilings remain specified as a percentage of M2 for auto-scaling, even though the money they move is the transactional subset.

12.3 Failure Mode 3 — Banking Liquidity

Tool 5 — Private Credit Guarantee Facility. Government guarantees overnight interbank reserve lending only — no corporate bonds, no mortgage-backed securities, no industry-specific support. Ceiling: 2.2 percent of current M2 (\$500B at launch). Closes 18 months post-trigger; no renewal.

Tool 6 — Rule-Based International Swap Lines. Pre-negotiated treaty amounts activate automatically. Treaty terms, eligible countries, and amounts are fixed at ratification — no discretion at activation. Ceiling: 1.3 percent of current M2 (\$300B at launch); auto-expiry 12 months per treaty terms.

12.4 Failure Mode 4 — Credit Cycle

Tool 12 — Equity Market Stability Reserve. Permanent fund of 1 percent of current M2 (\$227B at launch in Mode A; scales with M2 in Mode B, Mode C, and Mode Λ) buys the total market index only. Triggered when total market index falls more than 40 percent from 52-week high for 30 consecutive days, or when Constitutional Exits exceed 10 percent of Stable Floor pool over rolling 6-month window.

12.5 Cross-Mode Tools

Tool 3 — Velocity Support Mechanism. Transaction fees drop to zero; transfer limits lifted. No new money created.

Tool 13 — Naturalization Processing Resource Allocation. Additional federal fiscal resources allocated to clear the existing naturalization application backlog more rapidly during recessions. The monetary consequence is a byproduct of administrative efficiency, not the stated purpose.

12.6 Mode-Specific Tool Calibration

The seventeen tools are constitutionally available across Modes, but their utility, trigger calibration, and interaction with structural mechanisms vary by Mode; one, the Demurrage Backstop (Tool 16), is active only where the Mode targets deflation and lies dormant otherwise. The principle is consistency rather than override: each Mode's emergency capacity reflects the same proportional response to the same proportional shock, but the calibration must account for the structural mechanisms already operating within that Mode.

All emergency tool ceilings are specified as percentages of current M2 rather than fixed dollar amounts. This auto-scales emergency capacity across all Modes and over time — a tool sized to 3 percent of M2 provides the same proportional response whether M2 is \$22.4T or \$100T. The percentage-of-M2 specification strengthens rather than relaxes constitutional bounds.

Tools 1 and 8 (deflationary response tools) are largely redundant in Mode C and in Mode Λ during periods when KI is active, because both configurations provide automatic counter-deflationary expansion as a structural feature. Tools 1 and 8 remain available but their role is residual — they activate only if the structural KI expansion is insufficient. Mode A and Mode B rely on these tools more heavily because their structural mechanisms are more limited.

None of these Mode-specific calibrations create override authority. The FDCA cannot suspend constitutional bounds during crises. The seventeen tools remain bounded, rules-triggered, automatically reversing. Bounded but real beats unbounded but unaccountable.

12.7 Total Crisis Response Capacity

In a mature system, the combined first-year capacity of these seventeen tools is approximately 16 to 18 percent of current M2 in non-inflationary tools, plus up to 2 percent of current M2 in emergency K1 issuance. Tool 16 (the Demurrage Backstop) adds nothing to these totals: it is revenue-neutral by construction, redistributing its proceeds equally per capita, so it changes the incentive to hoard without creating or draining net money. At launch parameters the money-moving tools represent approximately \$3.7 to \$4.1 trillion in non-inflationary capacity plus \$447 billion in emergency K1. The increase over the prior fourteen-tool capacity reflects the addition of Tool 15 (M2 Contraction Floor, 3.0 percent of M2 annually), which addresses the Fisher debt-deflation cascade specifically. This compares favorably to the approximately \$2.9 trillion in genuine liquidity support the 2008 crisis required. The remaining \$5.9 trillion of the actual \$9.8 trillion deployed in 2008–2022 produced the 11.1 percent inflation of June 2022. The Citizens Standard provides sufficient capacity for genuine emergencies while being constitutionally incapable of the inflationary overreactions that current discretionary authority enables.

12.8 Emergency Tool Mapping for Custom Modes

Modes A, B, and C have explicit, tested tool calibrations documented in Section 14.6. A society that ratifies a custom Mode under Section 7.5 — including adaptive configurations like Mode Λ — inherits the seventeen-tool toolkit, but the toolkit's utility, trigger calibration, and interaction with that Mode's structural mechanisms may differ from the base-Mode calibrations in ways that are not always predictable in advance. One tool is conditional on the Mode's price-level target: Tool 16, the Demurrage Backstop, is active only for deflation-targeting configurations, so a custom Mode that targets mild deflation inherits it on the same terms as Mode A, while a price-stable or inflationary custom Mode carries it dormant, with no deflation-driven hoarding incentive

for it to offset. The constitutional ratification process for any new Mode must therefore include an explicit emergency tool mapping as a condition of FDCA verification.

The required mapping has four components. First, tool applicability: which of the seventeen tools apply without modification, which require Mode-specific trigger recalibration (as Tool 14 does across Modes A, B, and C), and which are structurally redundant given the new Mode's automatic stabilization mechanisms. Second, structural interaction: where the new Mode's formula-driven governors already provide automatic counter-cyclical response, the mapping must identify which tools become residual rather than primary — analogous to how Mode C's KI formula makes Tools 1 and 8 largely redundant. Third, gap identification: whether the new Mode's failure modes create scenarios that the existing seventeen tools cannot adequately address, and if so, whether an additional bounded tool should be ratified alongside the Mode. Fourth, ceiling scaling: confirmation that all tool ceilings expressed as percentages of current M2 remain proportionally appropriate given the new Mode's expected M2 growth trajectory.

This mapping requirement is not a barrier to Mode ratification — it is a transparency requirement. A society ratifying Mode Λ , for instance, would document that: the Demographic Governor and conditional KI make Tools 1 and 8 partially redundant in demographic-stress scenarios; Tool 14's trigger requires calibration relative to Mode Λ 's variable inflation target rather than a fixed target; and no structural gap exists in Mode Λ 's failure-mode coverage that the existing toolkit cannot address. The FDCA publishes the completed mapping alongside the ratification record, making it available for citizen review during the mandatory 90-day deliberation period. This ensures that when the Mode operates and a crisis arrives, citizens and the FDCA already understand which tools apply and how — rather than making those determinations under pressure.

13. The Market Exit

The framework's emergency tools, three-tier governance, statistical independence requirements, and constitutional protocol all serve to prevent corruption of the monetary system. They are designed to be sufficient. They may not be. Any sufficiently durable system must contemplate the possibility of its own failure and provide citizens with recourse when other safeguards prove inadequate.

13.1 What the Exit Provides

Under specific constitutional conditions, citizens may convert a portion of their Stable Floor balance into a non-domestic store of value of their choosing — outside the protocol, beyond the reach of any future protocol modification. The conversion is permanent for that portion. The Market Exit applies to Stable Floor holdings (K1 and K2 accounts) in all Modes. KI deposits are not subject to Market Exit because they are already unlocked and spendable.

13.2 Exit Destinations

The framework is intentionally agnostic about where the exit goes. Eligible exit destinations must meet five criteria: independence from US monetary policy, liquidity, international recognition, custody verifiability outside US-domiciled custodians, and auditable acquisition. Eligible destinations include physical gold and silver, foreign sovereign currencies of major trading partners, decentralized digital assets meeting liquidity thresholds, and FDCA-approved diversified foreign equity baskets.

13.3 Exit Conditions

Routine Exit: citizens above retirement age may exit up to 25 percent of their Stable Floor balance to any registered exit destination at any time. Time-locked over 5 years (1/5 per year).

Constitutional Exit: if specific constitutional violations occur, citizens of any age may exit any portion of their balance, including 100 percent. Qualifying violations include Tier 1 protocol breach, sustained inflation breach (CPI greater than 5 percent annualized for 24 consecutive months in Mode A or B; greater than 5 percent above target in Mode C or Mode Λ), identity system compromise, or foreign occupation/constitutional collapse. Time-locked over 18 months when triggered.

13.4 Strategic Effect

The Market Exit creates competitive pressure that benefits the framework even when not exercised. The protocol must remain attractive enough that citizens choose not to exit. The framework treats the citizen as the ultimate sovereign of the monetary system, not as a subject of it. The citizens are choosing to participate in this system; their continued participation is a vote of confidence — earned, not extracted.

14. Honest Limitations

A framework that does not acknowledge its failure modes is advocacy, not analysis. The limitations vary by Mode. This section presents them honestly, organized by which Modes they apply to.

14.1 Limitations Common to All Modes

Velocity collapse risk. $MV = PQ$. If velocity falls substantially, the price level falls regardless of M . Japan's velocity fell approximately 40 percent over 30 years. The Citizens Standard differs from Japan in eliminating the zombie bank problem through 100 percent reserve banking, but velocity decline from demographic aging is a structural force the emergency tools can only partially offset.

Shadow banking. The Citizens Standard has no jurisdiction over the \$20 trillion United States shadow banking sector or the \$13 trillion eurodollar market. Three partial mitigations apply within US jurisdiction: stablecoin issuers regulated as deposit

institutions, shadow credit instruments above \$1M registered on the public ledger, rehypothecation of collateral prohibited in regulated markets. International markets remain outside reach.

Identity verification infrastructure. The framework requires a national identity verification system that is comprehensive, fraud-resistant, privacy-preserving, and decentralized. Estonia operates such a system for 1.3 million people; the United States has 342 million. Deploying privacy-preserving cryptographic verification at US scale requires 3 to 5 years of infrastructure development and is a prerequisite for implementation.

Political durability. Every major institutional interest group loses something under the Citizens Standard: commercial banks lose the seigniorage value they currently capture from money creation (an implicit subsidy plausibly in the hundreds of billions of dollars a year), the federal government loses the ability to monetize deficits, the Federal Reserve loses its primary function. Major monetary reforms happen when crises delegitimize the existing system sufficiently that the coalition for change becomes larger than the coalition against it. The Citizens Standard cannot manufacture that crisis.

Multigenerational wealth concentration. Heritable Stable Floor balances introduce concentration dynamics that any equity-based capital stake system creates. The Citizens Standard's universal floor ensures even citizens inheriting nothing build a substantial capital position over their working life. Mode C's KI dividend further reduces dependence on inherited wealth by providing meaningful current income to every citizen.

14.2 Mode A-Specific Limitations

Wage rigidity. Wages are sticky downward. Under deflation, real wages rise automatically — beneficial for workers, but creates incentives to hire less and automate more. Workers without sufficient bargaining power to negotiate indexed contracts bear the adjustment cost.

Universal price indexing infrastructure. Mode A requires that all debt contracts be denominated in price-indexed units to prevent debt-deflation dynamics. Chile has operated this system since 1967, but deploying it across the US legal and financial infrastructure is a substantial implementation project. Mode B and Mode C do not require this.

14.3 Mode B-Specific Limitations

Loss of structural purchasing power gain. Mode B eliminates the rising-real-wages property that Mode A produces. Citizens must invest to preserve purchasing power; cash holdings do not appreciate.

Calibration sensitivity. Mode B's price stability target is more sensitive to calibration error than Mode A's deflation target. Mode A's slow indirect circulating expansion produces deflation across a wide range of K1 and K2 magnitudes. Mode B's calibration

must precisely match issuance to real growth; calibration errors produce inflation drift in either direction.

14.4 Mode C-Specific Limitations

Larger annual circulating-pool issuance. Mode C creates new circulating-pool money at a flat rate of about 1.98 percent of M2 (3.85 percent of the transactional circuit $M^{\wedge}T$) — approximately \$442 billion at launch — which scales gently with the economy thereafter as a constant share of a growing transactional circuit. This is the magnitude required to hit a 2 percent inflation target. A future society that finds approximately 2.3 percent of M2 (4 percent of the transactional circuit) in total steady-state annual issuance objectionable can amend to Mode A or Mode B by supermajority.

Political pressure on KI magnitude. A KI commitment of roughly \$442 billion at launch and scaling with the economy thereafter, paid as direct citizen distribution, is the largest single fiscal-equivalent flow in the framework. Institutional interests will pressure to capture, redirect, or restructure it. Mode C is meaningfully more politically vulnerable than Mode A or Mode B because more value is at stake in any successful capture.

Inflation target precision. The KI formula uses lagged data (trailing 12-month CPI, smoothed growth). Actual inflation will deviate from the 2 percent target year-to-year, typically within plus or minus 1 percentage point.

14.5 Mode Λ -Specific Limitations

Governor complexity. Mode Λ 's adaptive multipliers and boosters are formula-derived, but they are substantially more complex to verify and communicate than the base Modes' simpler parameters. A society that values constitutional simplicity and ease of citizen verification may find the base Modes more appropriate even if Mode Λ produces superior outcomes under demographic stress.

Multiplier calibration risk. The demographic and productivity governors require accurate, unmanipulated data from BEA, BLS, and Census. The framework's statistical independence commitment (Section 15.5) applies with particular force to Mode Λ , because incorrect inputs can activate multipliers inappropriately in ways that compound over decades. A CPI manipulation that would distort KI by one year in Mode C could distort a demographic multiplier in Mode Λ by fifteen years if the error persists.

Interaction effects. When multiple governors activate simultaneously — demographic stress plus productivity boom plus KI activation — the combined issuance can approach Mode C's circulating-pool expansion rates without having been ratified as Mode C. The constitutional guardrails in Section 9 (combined issuance cap of 3.5 percent of M2, KI 36-month sunset, multiplier review trigger) are designed to prevent this outcome, but they require active FDCA monitoring and citizen engagement to function as intended.

KI creep risk. The conditional KI stream can become quasi-permanent if the triggering conditions — particularly demographic stress — persist for extended periods. A society experiencing sustained population decline could find its Mode Λ functioning as a de facto

Mode C through repeated KI activations, without having made that constitutional choice explicitly. The 36-month KI sunset with mandatory reaffirmation vote is the structural defense; citizens must actively choose to continue each activation, preventing silent drift toward a different effective Mode.

Governance overhead. The mandatory FDCA review triggered at 15 consecutive years of elevated multiplier, the KI reaffirmation votes, and the annual multiplier publication requirements impose substantially more governance overhead than the base Modes. Societies that lack strong FDCA institutional capacity may find this overhead a practical obstacle to Mode Λ operation even if the formula itself is sound.

14.6 Limitations Not Solved by Mode Choice

Three limitations apply identically across all Modes. First, no countercyclical fiscal capacity: the balanced budget requirement eliminates Keynesian deficit spending; emergency tools provide approximately \$3.0 to \$3.5 trillion in first-year crisis capacity at launch parameters (rising to \$3.7–\$4.1 trillion with Tool 15), sufficient for genuine emergencies but cannot replicate the \$9.8 trillion deployed across 2008–2022. Second, equity market return dependence: the Stable Floor architecture depends on equity markets producing historically consistent returns; sequence-of-returns risk affects all Modes equally. Third, transition period vulnerability: the transition involves decades of parallel operation during which the system has the vulnerabilities of both architectures and the full protections of neither.

14.7 The Equity Transmission Architecture as Structural Feature

Several of the limitations above share a common cause: the framework deliberately routes monetary creation through capital markets rather than through banking. This is not an accident or a hidden vulnerability. It is the framework's most consequential structural choice.

Routing monetary creation through equity markets accomplishes four things that bank-credit transmission cannot. First, it eliminates internal Cantillon distortion: every citizen receives the same equity stake at issuance, with no early recipients capturing inflationary advantage. Second, it converts conventional seigniorage — currently a hidden subsidy to financial institutions — into citizen seigniorage: visible ownership of the productive economy by ordinary citizens. Third, it produces a structural store of real wealth — the Stable Floor — that survives nominal currency fluctuations. Fourth, it slows monetary transmission compared to bank-credit creation, which is the mechanism by which Mode A and Mode B avoid generating consumer-price inflation despite ongoing money creation.

Three costs are direct consequences of the equity-transmission architecture. First, sequence-of-returns risk: a citizen who turns 65 the year after a major equity-market crash holds a smaller Stable Floor than a citizen whose accumulation window ends after a long bull market. Under bootstrap resampling of the historical joint distribution, a substantial share of plausible histories — quantified in the companion empirical paper

(Neo-Solon, 2026b) on the same realizable basis, and larger for the earliest cohorts than the price-taker reconstruction implied — produce a locked Stable Floor below the median actual benchmark, a dispersion the standing K3 dividend and the framework's structural tail defenses are designed to offset. Second, broader market exposure: the framework's success is structurally tied to the long-run performance of the total US equity market. Third, foreign-capital interaction: foreign purchases and sales of US equities affect the Stable Floor's value just as they affect any other US-listed equity holdings.

A fourth consideration deserves explicit treatment: whether the Citizens Standard's equity-transmission architecture would itself alter long-run equity returns. This paper's projections do not assume the historical return continues unchanged: as Section 6.7 of the macroeconomic model (Neo-Solon, 2026e) establishes, the floor earns the general-equilibrium realizable return on capital, which falls below the historical price-taker average as the citizen capital share deepens. The intuition behind that result has two parts. First, the current system's approximately 6.5 percent annual M2 expansion has itself been a significant driver of historical equity valuations — asset prices have benefited substantially from the monetary expansion the Citizens Standard is designed to curtail. On this reasoning, Citizens Standard equity returns might run somewhat below historical, particularly in Mode A and Mode B where total new money entering through equity markets is only 0.35 to 1.2 percent of M2 annually, well below the current system's expansion rate. Second, the countervailing argument is that new money is directionally neutral for equity returns regardless of its source — whether it originates from bank credit or from citizen equity purchases, it represents demand for productive assets, and the Citizens Standard's K2 channel provides that demand directly and consistently rather than through the credit cycle. Section 6.7 makes the net effect precise: the floor earns the attenuated marginal product at its Mode's capture rather than the historical price-taker return, and the empirical paper (Neo-Solon, 2026b) characterizes the residual sensitivity through bootstrap resampling of the full historical distribution. The macroeconomic model (Neo-Solon, 2026e, Proposition 2 and Section 6.6) formalizes this question: it shows that the structural equity buyer produces a bounded, one-time valuation premium rather than a secular return collapse, because it is a net buyer only while the floors accumulate and becomes a bounded net seller as they mature — a deceased citizen's residual passes to an heir's floor rather than being sold — and because firm-side equity supply responds to valuation. This paper's projections use the general-equilibrium realizable return by Mode rather than the historical price-taker average, and the framework's structural advantages — universal participation, equal distribution, constitutional locking — are preserved across the full plausible return distribution regardless of its mean.

The honest case for equity-transmission architecture is comparative. The discretionary system makes commercial banks the primary transmission mechanism; the structural relationship between citizen and money is borrower-and-lender. The Citizens Standard makes equity markets the primary transmission mechanism; the structural relationship between citizen and money is owner-and-stake. Both are defensible. Both have failure

modes. A society choosing between these architectures is choosing which failure modes it prefers to live with — not whether to have failure modes.

15. Governance

15.1 The Constitutional Protocol

All monetary rules under the Citizens Standard are codified in a constitutional protocol — the supreme governing instrument of the system. The Federal Digital Currency Authority (FDCA) inherits the Federal Reserve’s operational functions but possesses zero discretionary monetary authority. It cannot set rates, create money, or take any action not explicitly authorized in the constitutional protocol.

15.2 The Three-Tier Updatability Framework

Tier 1 — Constitutionally Fixed

These provisions are unamendable. They are the load-bearing walls of the architecture. Tier 1 includes: the Model’s seven properties listed in Section 3.1; the prohibition on discretionary monetary creation outside the K1, K2, K3, KI, and emergency-tool channels; the constitutional lock on K1 and K2 Stable Floor accounts until age 65; equal per-citizen distribution at issuance for all channels; the structural requirement of the Composite Productivity Index, and that it and the price-level input be published through a tamper-evident, independently reproducible measurement pipeline; the payment-credit separation in banking; the four-failure-mode organization of emergency tools.

Tier 2 — Constitutionally Amendable (67 percent supermajority)

Parameters changeable through supermajority citizen vote with mandatory ninety-day deliberation. Tier 2 includes: the K1 base rate and its constitutional ceiling/floor; the K2 capture rate; the KI inflation target; the K2 calibration method; the Adaptive Smoothing Rule thresholds; bank leverage ratios; emergency tool ceilings, triggers, and reversal conditions; the choice of Mode; and — critically — the ratification of any new Mode beyond those already in force. This encompasses both the base Modes (A, B, C, D) and any adaptive or extended configurations such as Mode Λ. Any Mode ratified under Tier 2 must be verified by the FDCA to preserve the Model’s load-bearing properties.

Tier 3 — Operationally Updatable

Implementation details updatable by FDCA with public notice. Tier 3 includes: the specific data sources within each Composite Productivity Index category; cryptographic implementation of identity verification; the specific implementation of the tamper-evident measurement pipeline, whether a public append-only ledger or a reproducible hashed pipeline; technical infrastructure of distribution; FDCA reporting formats and publication schedules; and the specific banks authorized to provide transaction account custody services.

15.3 Mode Selection as Constitutional Choice

A society ratifies the Model — the Tier 1 architecture — and then chooses, as a Tier 2 question, which Mode to operate under. The choice is not technical optimization. It is a values question: does the society prefer rising real wages, nominal stability, visible citizen dividend, or an adaptive configuration that responds to demographic and productivity conditions?

Mode selection requires the same 67 percent supermajority and 90-day deliberation as any other Tier 2 amendment. A society can ratify the Model in one constitutional convention and select Mode in the same vote, or it can ratify the Model and defer Mode selection to a subsequent vote. A society that has operated under one Mode for decades can change to another Mode by supermajority — including to a Mode not presented in this paper, provided it preserves the Model's load-bearing properties. The architecture supports all these pathways and remains open to others.

Supplementary replication. That this is a values question rather than a technical optimization is borne out quantitatively. Treating the dividend share κ_d as a continuous dial and scoring each setting by social welfare — using the realizable return the architecture assigns each configuration (5.4 percent at the floor-only split, 4.3 percent at Mode B's 60/40), a liquidity-constrained population share from the Federal Reserve's SHED survey (37 percent of adults could not cover a \$400 emergency in cash), and the Model's own marginal-propensity-to-consume schedule — the welfare function is monotonic in κ_d and has no interior maximum: there is no welfare-optimal Mode to compute. What the dial trades is immediate liquidity (a larger spendable dividend) against compounding ownership (a larger locked stake earning a higher return), bounded above by the inflation ceiling, with neither side dominating. This is the formal basis for assigning Mode selection to a supermajority vote: the choice is genuinely constitutional rather than technical. A polity may use the liquidity-constrained share as a reference point for the spendable share and may favour a Mode whose dividend rises countercyclically, but these are inputs to the vote, not an optimum. Runnable as the mode-choice-welfare module in the distribution replication package.

This is a meaningful innovation in monetary governance. Current monetary policy is set by appointed committees with no formal citizen input. The Citizens Standard makes the most consequential monetary question — what inflation regime should we have, and how adaptive should our issuance formula be? — an explicit constitutional decision made by citizens directly.

15.4 Mode Stability and Why Modes Do Not Auto-Rotate

A natural question follows from the multi-modal architecture: do Modes rotate or oscillate inside the system itself, automatically, when conditions change? The answer is no — and the reasoning is load-bearing for what kind of framework the Citizens Standard actually is.

Mode selection is exclusively a Tier 2 constitutional choice. Changing from one Mode to another requires the same 67 percent supermajority and 90-day deliberation as any

other Tier 2 amendment. The framework cannot change its own Mode automatically based on inflation breaches, recessions, demographic shifts, or any other condition. Auto-rotation would violate the framework's central commitment in three structurally important ways.

First, auto-rotation would reintroduce discretion through the back door. If the framework can decide to shift from Mode A to Mode C based on some metric, the question becomes: who chose that metric, and who decided what threshold triggers the shift? That is committee discretion at the rule-design layer rather than the operational layer. The Citizens Standard rejects discretion at both layers.

Second, auto-rotation would collapse the meaning of Mode selection. The whole point of letting citizens choose Mode is that the choice is consequential and durable. A Mode that auto-rotates is not really a choice; it is a default.

Third, auto-rotation would weaken the Market Exit's function. The Market Exit is the citizen's structural recourse when constitutional bounds are breached. If Mode auto-rotates, the constitutional promise is moving, and the Exit trigger becomes ambiguous.

What does adapt automatically within a Mode: K1 calibration scales with GDP per capita. K2 calibration scales with real growth. K1 calibration in Mode C scales with the cumulative inflation gap by formula. In Mode A, the governors respond to demographic and productivity data by formula. The bank leverage ratio tightens and loosens countercyclically. Emergency tools activate by formula triggers and sunset automatically. The Long-Run Stability Provision in Mode A automatically cycles parameters toward Mode B equivalence as the GDP/M2 ratio crosses thresholds. All of this is adaptation within a ratified Mode — not rotation across Modes. The principle remains intact: the rules adapt; the choice of rules does not.

The principle that emerges from this analysis is straightforward: the framework adapts within constitutional bounds; constitutional bounds adapt only by citizen action. Anything that can be calculated by formula from observable economic variables is automatic. Anything that requires choosing what to value — purchasing power preservation, nominal stability, visible citizen dividend, adaptive demographic response — is a citizen decision.

15.5 Statistical Independence

Under the Citizens Standard, the statistical agencies that measure GDP, the price level, velocity, and unemployment become the effective central bank. This is a structural consequence of replacing committee discretion with formula-based issuance. The agencies do not set policy; they measure the inputs that determine policy. But they do set the inputs, and the framework is therefore only as good as the inputs are.

The required defenses are constitutional rather than legislative. Statistical agencies must possess independence equivalent to what the Federal Reserve Board holds today: fixed multi-year terms for agency directors with removal-for-cause only; operational budgets that cannot be cut by executive action below specified floors; methodologies publicly

documented in advance, changeable only through transparent supermajority process with mandatory deliberation periods; raw data inputs preserved and independently auditable; foreign-government replication of key indices as cross-checks against domestic manipulation.

This requirement is most acute in Mode C, where the KI formula is directly driven by published CPI on annual cadence, and in Mode Λ , where demographic data drives multipliers that compound over decades. It is least acute but still present in Mode A, where CPI affects K1 calibration through the price-level adjustment formula but does not drive the dominant K1+K2 magnitudes. The framework's overall claim — that rules are better than discretion — depends entirely on the rules' inputs being honest.

15.6 Identity Verification

The protocol requires a decentralized identity system that verifies unique, living personhood — answering whether a person exists and is counted exactly once, without requiring disclosure of any personal information beyond that binary fact. Privacy-preserving cryptographic methods, including zero-knowledge proofs, ensure that verification does not create a surveillance infrastructure. The identity system is the prerequisite for K1 (verifying citizenship events), K2 (counting living citizens for distribution), and KI (preventing duplicate dividend claims).

16. System Comparisons

The Citizens Standard is not a new idea. Each of its components has been proposed independently by serious economists. The Citizens Standard's contribution is structural integration where each component is load-bearing for the others, and the introduction of the Mode-selection architecture that allows one Model to host multiple coherent monetary regimes — from simple single-governor configurations to complex adaptive ones.

16.1 Comparison with Current System

The current system creates money through bank lending and Federal Reserve discretion. Money supply has grown at approximately 6.5 percent per year long-run. Distribution is hierarchical (banks first). Discretion is maximum. The Citizens Standard creates money at 0.35 percent of M2 annually (Mode A), 2.0 percent (Mode B), or 2.3 percent (Mode C), distributes equally to citizens at issuance, and operates under constitutional rules. Inflation protection is fundamentally different across all Modes compared to the current system's 96 percent purchasing power loss over 113 years.

16.2 Comparison with MMT

Modern Monetary Theory observes that a sovereign currency issuer faces no operational financing constraint and proposes fiscal policy as the primary tool of macroeconomic management. The Citizens Standard agrees with the empirical

observation and rejects the normative conclusion. The technical capability to create unlimited dollars should be constitutionally constrained, not exercised. Mode C, the closest of the base Modes to MMT in monetary expansion magnitude, differs from MMT decisively in distribution: KI goes equally to citizens, not to government programs determined by legislative process.

16.3 Comparison with Bitcoin

Bitcoin shares with the Citizens Standard a commitment to constraining monetary expansion through structural rules. The frameworks differ on three dimensions. First, distribution: Bitcoin's new issuance flows to miners; the Citizens Standard's K1, K2, K3, and KI distribute new money to citizens. Second, scale calibration: Bitcoin's issuance is unconnected to real economic conditions; K2 is explicitly calibrated to real growth, KI to inflation gap. Third, structural floor and dividend functions: Bitcoin provides no Stable Floor or citizen dividend; the Citizens Standard provides both.

16.4 Comparison with UBI Proposals

Mode C is the configuration most directly comparable to Universal Basic Income proposals. KI is funded through citizen seigniorage rather than taxation; KI is calibrated by formula rather than legislatively; KI is constitutionally protected rather than subject to ordinary legislative reversal. UBI proposals and Mode C agree that direct equal-distribution citizen payments are valuable; they disagree about the appropriate funding source and constitutional status of those payments.

16.5 The Citizens Standard's Distinctive Position

Across these comparisons, the Citizens Standard occupies a position none of the alternatives reach. It accepts MMT's observation that the dollar is a public utility but rejects political control over its creation. It accepts the Chicago Plan's separation of money and credit and adds the missing piece (citizen-level distribution of citizen seigniorage). It accepts Bitcoin's scarcity insight and adds the distribution mechanism Bitcoin lacks. It accepts UBI's distribution insight and adds the funding mechanism UBI lacks. The Mode-selection architecture — from simple base Modes to adaptive configurations like Mode Λ — allows a society to choose which of these adjacent frameworks the Citizens Standard most resembles in operation while preserving the same underlying Model.

17. International Applicability

This paper is written for the United States. The Citizens Standard's core architecture is adaptable to any sovereign state, but not every state equally.

17.1 Requirements for Any State

Implementation requires six conditions: monetary sovereignty over a domestic currency; constitutional or legal capacity to bind future governments to monetary rules; population verification infrastructure; a locked investment vehicle (for Modes targeting Stable Floor accumulation); minimum GDP for meaningful per-citizen K2 amounts; and financial inclusion sufficient to give every citizen a Stable Floor account.

17.2 Mode Choice and National Context

Different national contexts may favor different Modes. A small economy with high savings rates and aging demographics may prefer Mode A's purchasing power preservation, or an adaptive Mode like Mode Λ that adjusts for demographic stress. A large economy with extensive long-duration nominal contracts may prefer Mode B's stability. A developing economy with significant unbanked or low-income population may prefer Mode C's visible citizen dividend. The Mode architecture allows national adaptation without requiring departures from the Model.

17.3 Reserve Currency Implications

The United States dollar functions as the world's primary reserve currency. The Citizens Standard's structural design strengthens the monetary case for the dollar in any Mode — even Mode C produces a 2 percent inflation rate that is more credible and predictable than the discretionary system's average. Four of the five inputs to the Composite Productivity Index are independently auditable by foreign governments. No current reserve currency offers this verifiability. The framework does not require international coordination to function.

17.4 The Adaptable Core

Stripping away US-specific parameters, the Citizens Standard's adaptable core is the Model's seven properties listed in Section 3.1. These can be implemented across a wide range of national contexts with any Mode parameterization. Policymakers evaluating implementation in non-US contexts should note that all projected Stable Floor values in this paper are calibrated to US equity market return history. Countries with lower historical long-run equity returns would produce proportionally lower absolute Stable Floor outcomes while preserving the framework's structural participation advantages. The democratic accountability requirement — that rules cannot be changed without supermajority citizen approval — is non-negotiable in any honest adaptation.

18. Long-Run Considerations

The framework as specified is designed for indefinite operation. This section addresses three categories of long-run dynamics: intra-Mode A temporal cycling, inter-Mode constitutional change, and structural risks that apply across all Modes.

18.1 Intra-Mode A Cycling

In Mode A, frozen circulating M2 against compounding real growth produces accelerating deflationary pressure. The GDP/M2 ratio rises monotonically as long as real growth exceeds the equity-seller-channel recirculation rate of $K1+K2$ dollars. Numerical simulation under base-case 2 percent real growth shows the GDP/M2 ratio reaching 5× the launch ratio between approximately year 80 and year 120. The Long-Run Stability Provision automatically activates a recalibrated mode (functionally equivalent to Mode B) when the GDP/M2 ratio exceeds 5.0× the implementation-date ratio for two consecutive years, and returns to standard Mode A operation when the ratio falls below 4.0×. This provides a default pathway from Mode A to Mode B equivalence in societies that do not amend Mode by explicit Tier 2 vote.

18.2 Inter-Mode Constitutional Change

A society can change Mode by Tier 2 amendment. Mode A to Mode B is the natural transition path as the economy matures and long-duration nominal contracts proliferate. Mode A to Mode C is the path for societies that decide a citizen dividend is constitutionally desirable. A society could also transition to Mode Λ or another adaptive configuration if it concludes that formula-driven demographic response is constitutionally preferable to relying solely on emergency tools. All reverse transitions are equally available. Mode changes are deliberate, infrequent, constitutional events — possibly occurring once per several generations, possibly never.

18.3 Structural Risks Across All Modes

Three categories of long-run risk apply across all Modes. First, statistical agency capture: if the agencies producing GDP, CPI, or population data are politically captured, the K formulas produce wrong outputs even with formula integrity preserved. This risk is heightened in Mode Λ , where demographic data drives multipliers that compound over decades. Second, identity system compromise: if the identity verification system is compromised at scale, fraudulent citizenship claims could divert $K1$ and KI issuance. Third, evolution of the broader economy: GDP, M2, and CPI are meaningful measures today; they may become less meaningful in an economy dominated by digital goods or AI-generated value. The Tier 2 amendment process allows future generations to adapt the formulas.

A fourth structural consideration is labor displacement from automation and robotics. In labor-anchored monetary systems, productivity gains from automation accrue primarily to capital owners while simultaneously compressing the wage-based distribution channel — a dynamic that tends to widen inequality as automation accelerates. The Citizens Standard's distribution mechanism is anchored to population and productivity rather than to labor income. Rising automation increases the Composite Productivity Index, which increases $K2$ calibration, which increases Stable Floor deposits for every citizen — inverting the distributional dynamic of debt-based systems. Separately, the dual-circuit architecture's equity-anchored household balance sheets mean that automation-driven

deflationary pressure does not trigger the debt-deflation spirals that characterize labor-anchored systems under similar conditions. These structural properties do not eliminate the social disruption of large-scale labor displacement; they do mean the monetary architecture does not amplify it.

18.4 The Constitution Principle

The framework follows the principle of the United States Constitution: provide robust architecture for the foreseeable horizon, with explicit mechanisms allowing future generations to adapt the framework as conditions evolve. The Citizens Standard does not attempt to solve year 2200 problems. It attempts to provide architecture that year 2200 citizens can continue using or amend by supermajority. Both outcomes are acceptable; both are improvements over discretionary committee judgment.

18.5 A Complementary Fiscal Option: Land Value Taxation

The framework is deliberately agnostic about taxation, which remains a sovereign political choice; nothing in the monetary design prescribes a revenue system. One fiscal instrument is nonetheless unusually aligned with the framework's own logic, and is noted here as a complementary option a society may choose to pair with the architecture — not a requirement. A land value tax falls on the unimproved value of land — an asset that is not produced, and whose appreciation therefore accrues to owners by position rather than by contribution. This is the real-economy mirror of the Cantillon distortion the framework removes from money creation: just as the architecture refuses to let proximity to issuance capture unearned monetary advantage, a land value tax refuses to let proximity to location capture unearned economic rent. Three properties make it complementary rather than merely compatible. It is non-distortionary, falling on rent rather than on the labor or capital whose productivity the issuance channels are calibrated to track. It damps land and housing price growth at its source — a structural driver of measured cost-of-living that monetary contraction can only address bluntly — on a base the measurement-integrity layer can observe and that is difficult to conceal or relocate. And it supplies a durable, evasion-resistant revenue foundation that speaks to one of the framework's acknowledged limitations, the absence of countercyclical fiscal capacity: not by providing deficit capacity, but by giving a society a stable fiscal base it may choose to build on. The framework neither requires nor assumes it; the choice remains with politics. It is noted here only because the synergy is specific rather than generic.

18.6 Payout Neutrality and the Equity-Transmission Safeguard

The equity-transmission mechanism directs the K1+K2 deposit flow into a permanent, schedule-driven purchase of the total-market equity index — operationally the CRSP US Total Market Index (the basis of the largest total-market index funds, e.g. Vanguard's VTI), or equivalently the Wilshire 5000 Total Market Index, a float-weighted index spanning essentially all US-listed equities. For that flow to fund real capital formation rather than to bid up the price of a fixed share supply, the aggregate net supply of equity

must respond positively to valuation: in the presence of the structural buyer's premium, firms must on net issue equity, not retire it. The prevailing regime runs the other way. For domestic nonfinancial corporations, net equity issuance — gross issuance less retirements through repurchases and cash mergers — has been negative since the mid-1990s (Federal Reserve, Enhanced Financial Accounts; SEC 2020), and repurchases are procyclical: firms buy back disproportionately when prices are high and withdraw when they are low (Baker and Wurgler 2002; Jiang and Koller 2011). Over 2010–2019, net repurchases averaged about 1.5% of market capitalization per year (BIS 2020).

That figure is the crux. The structural-buyer flow into equity — $K1$ plus $K2$, the locked floor (funded at 60 percent of the post-citizenship real-growth budget under Mode B's 60/40 split, with the remaining 40 percent paid out as the $K3$ dividend) — is approximately 0.39% of US equity market capitalization per year (Neo-Solon, 2026b; 2026c). In general this flow is $A^* = (1 - \kappa_d) \times g_r \times M2 + \kappa_d \times K1_{agg}$, or as a share of the domestic index capitalization M_{index} , $f = A^*/M_{index} \approx (1 - \kappa_d) \times g_r \times (M2 / M_{index})$, the approximation holding because the citizenship endowment $K1_{agg}$ is small relative to the growth bulk. Any adopting state recovers its own structural-buyer flow from its money supply, its index capitalization, and its chosen dividend share κ_d — the base-Mode value $\kappa_d = 0$ giving 0.65 percent and the Mode B value $\kappa_d = 0.4$ giving 0.39 percent. The corporate sector has therefore been withdrawing float on net at roughly twice the rate the structural buyer would add demand. Absent a constraint, the buyer's purchases fall on a contracting share supply, and the mechanism finances repricing instead of investment — defeating its own purpose.

The Citizens Standard therefore establishes, at constitutional level, a payout-neutrality safeguard: corporate payout policy may not be permitted to neutralize the equity-transmission mechanism. The architecture mandates a constraint on net share repurchases sufficient to keep aggregate net issuance positive under the structural-buyer premium. Its operative form — a graduated repurchase charge with a hard ceiling — is specified by statute (Neo-Solon, 2026d, §12); its necessity and calibration target are established by the structural-buyer analysis (Neo-Solon, 2026h).

Three properties make this safeguard consistent with the framework's governance commitments and ensure it concentrates no new power:

It is mechanical and locked, not discretionary. The safeguard is a published, fixed schedule, alterable only through the Tier-2 constitutional amendment process. No official may set, waive, suspend, or vary it case-by-case. This is the same anti-discretion property that governs the index definition (Neo-Solon, 2026h, §9) and the constitutional locks (§15): the rule, not an administrator, decides.

It is administered outside the monetary authority. The charge operates through general corporate tax law and any conforming securities-law conditions through the existing securities regulator. The FDCA — the entity that executes equity purchases for the floors — has no role in setting or enforcing it. The safeguard therefore adds no power to the FDCA, which remains a price-insensitive, schedule-bound buyer; and by

the mirror-voting rule (Neo-Solon, 2026h, §11) the FDCA's equity stake conveys no corporate control in any case.

It is payout-neutral, not payout-suppressing. The safeguard constrains the form of distribution (net repurchases), not the right of shareholders to a return. Dividends remain unconstrained, so capital is still returned to owners — including, through their Stable Floors, the citizens themselves. The intent is to redirect distribution toward issuance-and-investment and toward dividends, not to trap earnings. A documented side effect is lower corporate leverage, since buybacks have served as a leverage-targeting tool and reinforce the effect of debt issuance on capital structure (BIS 2020) — a financial-stability gain, not a cost, for the system citizens depend on.

The safeguard is thus a constraint on the corporate sector operated through existing legal channels, locked against discretion, and neutral-to-beneficial with respect to citizens' returns — a complement the equity-transmission mechanism requires, not an expansion of any institution's power.

19. Situating the Citizens Standard in the Literature

The Citizens Standard is not a single reform but a synthesis drawn from three largely separate bodies of work — the rules-versus-discretion debate over monetary authority, the full-reserve banking tradition, and the broad-capital-ownership and social-dividend tradition — together with definite positions in the free-banking and Modern Monetary Theory debates. This section situates the framework in each, states what it borrows and what it does not, and identifies the gap it is meant to fill. The claim to originality is not that any single component is new; most are old. It is that the particular combination — a non-discretionary constitutional issuance rule, distributed equally at creation, as locked individual equity in the productive economy, atop full-reserve payment-credit separation — has not previously been assembled.

19.1 Rules versus Discretion, and the Constitutional-Money Tradition

The case against discretionary monetary authority is long established. Friedman (1960) argued that a central bank could not reliably out-perform a simple constant rule for money-supply growth — the k-percent rule — given long and variable policy lags. Kydland and Prescott (1977) gave the argument its modern form: because a discretionary authority cannot credibly commit, optimal-seeming discretion is time-inconsistent and yields worse outcomes than a binding rule. Buchanan, in the constitutional-political-economy tradition, pushed further — monetary arrangements should be constitutional choices rather than ongoing policy decisions (Buchanan 1962; Brennan and Buchanan 1981) — and Hayek (1976) reached a kindred distrust of monopoly issuance from the opposite direction, proposing competing private currencies. Against these, dominant contemporary practice is discretion with a feedback rule — the Taylor rule (Taylor 1993) and its market-monetarist cousin, nominal-GDP targeting (Sumner 2021) — both of which retain a discretionary authority and reform only its objective.

The Citizens Standard belongs to the rules-and-constitution side of this divide, closest in spirit to Buchanan's monetary constitution and Friedman's binding rule. It departs from both in two respects: where Friedman fixes a single number, the framework anchors issuance to a smoothed multi-source productivity index, so the rule tracks real output without becoming discretionary; and where most rule proposals remain ordinary policy a future authority can revise, the framework entrenches the rule constitutionally and treats the choice among regimes as a supermajority decision.

19.2 Free Banking: The Decentralized Alternative

The most fully developed rival remedy is free banking (Selgin 1988; White 1984; Selgin and White 1994), which holds that money should be created competitively by many private banks and disciplined through convertibility rather than by any central authority, reading the Scottish experience as evidence that decentralized issuance can be stable. The Citizens Standard shares the diagnosis — discretionary central authority is the problem — but adopts the opposite remedy, and the contrast is worth stating precisely.

Free banking decentralizes issuance and disciplines it through convertibility; the Citizens Standard centralizes the issuance rule while leaving credit fully decentralized and competitive, so banks continue to assess risk and allocate capital as a free-banking theorist would want. The efficiency argument for decentralization therefore applies to credit allocation, which the framework preserves, not to the growth of the base, where slow and predictable behavior is the objective. The genuine disagreement is whether convertibility-disciplined competitive issuance is a stable, durable equilibrium; the historical record is equivocal — antebellum US note discounts rose with distance largely independent of issuer quality, a textbook informational-asymmetry failure (Goodhart 1988) — and free banking was everywhere abandoned and nowhere re-established, which is itself evidence about its durability.

19.3 Full-Reserve Banking and Payment-Credit Separation

The framework's banking architecture descends from the full-reserve, or 100-percent-money, tradition. Fisher (1935) proposed 100 percent reserves against checking deposits, severing money creation from lending; the idea was central to the Chicago Plan associated with Simons (1936), and Friedman (1948) endorsed it in his early work. Tobin (1985) advanced a related deposited-currency proposal. The tradition was revived analytically by Benes and Kumhof (2012), whose evaluation found that a Chicago-Plan transition could reduce debt and dampen the credit cycle, and politically by Positive Money (Jackson and Dyson 2012), the Icelandic monetary-reform report (Sigurjónsson 2015), and the 2018 Swiss Vollgeld referendum. The Citizens Standard adopts the load-bearing feature — full-reserve, constitutionally protected payment accounts; at-risk term deposits; a payment system insulated from the credit cycle. Its departure is in what happens to the new money: the 100-percent-money literature almost uniformly routes the seigniorage to the government, through spending, tax reduction, or debt cancellation, and treats distribution as secondary, whereas the Citizens Standard distributes the issuance equally to citizens as locked individual equity. The banking reform is borrowed; the distribution is the novel element. The framework should also concede, rather than minimize, the standing critique of narrow banking — the boundary problem (Goodhart 2008): highly liquid term claims tend to migrate into transactional use, so the separation is never as clean in practice as on paper. The honest claim is partial — full reserves raise the cost and visibility of near-money creation without abolishing it.

19.4 Broad Capital Ownership and the Social Dividend

The framework's distributional layer connects to a tradition usually kept separate from monetary reform. Meade (1989) proposed a social dividend funded by publicly held productive capital; the Alaska Permanent Fund has paid a resource-financed citizen dividend since 1982, and Norway's Government Pension Fund Global demonstrates large-scale public equity stewardship. More recent proposals include Bruenig's (2018) collectively owned social wealth fund, Hamilton and Darity's (2010) baby bonds, the broad-ownership program of Blasi, Freeman, and Kruse (2013), Farmer's (2016) proposal for state trading of a broad equity index, and Piketty's (2020) universal capital

endowment financed by progressive taxation. A parallel cash-transfer strand runs from Friedman's negative income tax through Van Parijs (1995) to recent universal-basic-income proposals. The Stable Floor is distinguished on funding and form: it is financed neither by taxation nor by resource royalties but by monetary seigniorage routed to citizens at issuance, and it takes the form of an individually owned, constitutionally locked equity account compounding over a lifetime rather than a collective fund or a flat dividend. The framework's KI channel, when active, is a basic-income-like dividend, but in the framework it is secondary to the equity Stable Floor and is calibrated against an inflation target.

19.5 Modern Monetary Theory as Contrast

Modern Monetary Theory (Mosler 1996; Wray 2012; Kelton 2020) shares one premise with the framework — that money is fundamentally a creature of the state rather than of private banks — and reaches nearly opposite conclusions. MMT treats discretionary fiscal policy, anchored by a public job guarantee and constrained only by inflation, as the proper vehicle for the state's monetary capacity. The Citizens Standard agrees that issuance is a public prerogative but rejects discretion in favor of a non-discretionary formula, and routes new money to individually owned citizen equity rather than to government expenditure. The disagreement is not over whether the state can issue, but over whether it should do so at discretion and through the budget.

19.6 The Gap the Framework Occupies

Each component of the Citizens Standard has a respectable ancestor: the binding rule from Friedman and the constitutional framing from Buchanan; the full-reserve payment architecture from Fisher, Simons, and Benes and Kumhof; the broad-capital-ownership aim from Meade, Alaska, and Piketty. What the prior literature does not contain is the combination. No existing proposal joins a non-discretionary, constitutionally entrenched issuance rule; equal per-citizen distribution at the point of issuance; in the form of locked individual equity in the productive economy; atop full-reserve payment-credit separation; with the choice among inflation regimes left to citizens by supermajority. The rules tradition supplies the discipline but not the distribution; the full-reserve tradition supplies the banking architecture but routes the proceeds to the state; the social-dividend tradition supplies the distributional aim but funds it through taxation and typically as a collective fund or a cash flow. The framework's contribution is the synthesis, together with the mechanism that makes it coherent — seigniorage-funded individual equity, constitutionally bound.

20. Conclusion

The central claim of this paper is structural. Money should not be created by committee. The rules of monetary creation should be transparent, formula-based, and ratified by citizens — not opaque, judgment-based, and set by appointed officials. The proper

subject of constitutional politics is the inflation regime itself, not just the institutions that administer it.

The Citizens Standard provides architecture that makes this possible. The Model — dual circuits, citizen-anchored issuance, separated banking, constitutional governance — is the load-bearing structure. The four base Modes — Deflationary, Stable, Inflationary, and Distributed — are coherent monetary regimes that the architecture can host. Mode A demonstrates that the architecture can also host adaptive multi-governor configurations. A society ratifies the Model and then chooses the Mode that reflects its values.

The paper presents projected Stable Floor balances at age 65 because those figures are auditable and falsifiable — they give economists and policymakers something concrete to evaluate. But those numbers, while accurate, describe only the most easily measured consequence of the framework. The deeper consequence is institutional. The Stable Floor is not a retirement product. It is a constitutional guarantee that every citizen, from the moment of citizenship, holds a permanent and growing ownership position in the national productive economy, funded not by taxation but by citizen seigniorage — the value that monetary creation produces, routed to citizens at issuance rather than captured by institutions. No labor-market success is required to earn it. No behavioral discipline is required to maintain it. No inheritance is required to seed it. It exists because the citizen exists, and it grows because the economy grows.

This changes more than retirement outcomes. It changes the decision environment every citizen operates in across their entire working life. A citizen who knows from birth that absolute destitution in old age is constitutionally impossible bargains differently in labor markets, tolerates risk differently in career and entrepreneurship, resists predatory debt differently, and relates to civic institutions differently than one who does not. These second-order effects are not proven by this paper — they are the research questions the framework generates, and they are likely more consequential than the retirement-income comparisons in Section 7. The empirical paper (Neo-Solon, 2026b) documents what the framework produces in measurable terms. The companion transition paper (Neo-Solon, 2026c) specifies the migration path from the current system to that destination — the phases, timelines, and hard transition problems the framework must navigate to become operational. The larger question — what happens to a society in which every citizen possesses constitutionally guaranteed long-duration capital ownership from birth — belongs to the political economy literature this framework is intended to seed.

The framework is also not easily categorized by existing political vocabulary. It is not socialism — it does not equalize outcomes or eliminate private accumulation. It is not welfare liberalism — it requires no ongoing tax transfer and no means-testing. It is not traditional capitalism — it routes citizen seigniorage into universal citizen equity rather than into institutional balance sheets. It is closer to what political economists have called a property-owning democracy, or what might be called constitutional civic capitalism: a system in which market inequality is permitted and even celebrated, but in which the floor beneath every citizen is guaranteed not by redistribution but by architecture. The philosophical lineage runs through James Meade's property-owning democracy, social

dividend theory, and the distributist tradition — but the mechanism is monetary rather than fiscal, and the governance is constitutional rather than legislative.

The Modes presented here demonstrate the architecture's range — rising real wages, nominal stability, visible citizen dividend, demographic adaptiveness — but the framework supports any coherent Mode a society ratifies. None of these choices is correct or wrong in some external sense. Each is a constitutional choice properly made by citizens through supermajority process.

The current system makes none of these choices available. It produces approximately 6.5 percent annual M2 expansion, hierarchical distribution to financial institutions, and 96 percent purchasing power loss across a century. It does this not because any of these outcomes was deliberately chosen but because no mechanism exists for any choice to be made. Committee discretion is not a choice; it is the absence of choice. The Citizens Standard's most fundamental claim is that monetary outcomes are too consequential to be the residual byproduct of unaccountable institutional process.

Milton Friedman argued in 1960 that monetary policy should be governed by rules rather than discretion — that a computer executing a simple formula would outperform any committee over time. He was right. He did not have access to the distributed protocol technology, the verified digital identity infrastructure, or the investment vehicles that make a citizen-based monetary rule practical today. He also did not consider that the rule itself might be a constitutional choice from a defined menu, or that every citizen might hold, from birth, a constitutionally guaranteed equity stake in the economy the rule governs. The Citizens Standard is an attempt to build what Friedman described, with the addition that the rule is selected by citizens rather than imposed on them, and that the stake is theirs from the moment they join the system.

The framework presented here is not complete. The architectural balance among K1, K2, K3, and KI — and any future adjustment — requires constitutional choices this paper cannot make on behalf of the public. The transition sequencing requires detailed legislative drafting. The empirical backtesting of the Composite Productivity Index against historical periods would strengthen the framework's analytical foundation. The Mode-selection mechanism requires a constitutional design more detailed than this paper provides. These are acknowledged gaps, not hidden ones.

The invitation of this paper is not final adoption but serious engagement — from economists who will find errors in the analysis, from policymakers who will identify implementation obstacles, and from citizens who will recognize in this proposal something they have always known: that the money in their pocket should work for them, not against them, and that the rules governing it should be theirs to choose.

The pen name Neo-Solon was chosen deliberately. The Athenian reformer of 594 BC cancelled the debts that had reduced citizens to bondage, restructured the monetary and legal system to prevent recurrence, and understood that a republic cannot sustain democratic equality while economic bondage exists within it. This framework attempts the same dual reform: constitutional rules that prevent the monetary mechanisms of bondage from operating, and the direct recognition that every American who joins this

system — whether by birth or by naturalization — deserves a stake in it from the moment of citizenship. Not a pension. Not a transfer. A stake.

Glossary

Citizens Standard — A constitutional monetary protocol with dual-circuit architecture, citizen-anchored issuance, multi-channel issuance design, and constitutional governance. Hosts a constitutionally selectable menu of Mode configurations, from simple base Modes to adaptive multi-governor configurations.

Model — The unamendable architecture: dual circuits, citizen-anchored issuance, equal per-citizen distribution at issuance, equity-based ownership, separated banking, constitutional rules, supermajority amendability with citizen-only voting. Preserved across all Modes.

System — The specific Mode parameterization a society chooses by Tier 2 amendment.

Mode — A constitutional configuration of the framework: the specific menu choice a society ratifies. The four base Modes (A, B, C, D) differ in inflation outcome, KI activation, distribution form, and circulating-pool dynamics. Additional Modes — including the price-stability solver (Mode Ω), the zero-issuance corner (Mode 0), and adaptive multi-governor configurations (Mode Λ) — may be ratified by the same Tier 2 supermajority process provided they preserve the Model's load-bearing properties.

KT — The transition-only issuance channel. Creates money directed to Legacy Debt Trust bond redemption rather than to citizens, calibrated to a price-level path at approximately 1.5 percent of M2. Active only in Mode T; self-throttling on consumer inflation; throttles to dormancy once the public debt stabilizes within its operational band (~30 to 60 percent of GDP), thereafter available as a symmetric open-market instrument. Distinct from KI, which is a citizen-dividend channel.

Mode T — The transition configuration. Runs K1 and full-rate K2 — sixty percent to citizen Stable Floors, forty percent as the standing K3 dividend — at approximately stable prices (near a one-half transaction ratio) while the KT channel retires pre-existing public debt. Self-limiting: as public debt enters its operational band (~30 to 60 percent of GDP), KT throttles to dormancy and the system continues as Mode T-stable, with KT retained as a symmetric open-market instrument.

Mode T-stable — The automatic permanent steady state that follows Mode T once public debt is retired; it is Mode Ω . K1 and full-rate K2 continue funding Stable Floors, the split is solved for zero derived inflation, and KI is available to correct drift in either direction (KT is inactive). At the one-half transaction ratio this coincides with the fixed Mode B split; for any other economy it is the Ω solve. Requires no vote; a society may optionally adopt a different steady-state Mode (A-like, C-like, or Λ) by Tier 2 supermajority.

Mode A — Deflationary System — Targets approximately 1.9 percent annual deflation. K1 and K2 active, KI inactive. Produces structurally rising real wages and a modest real Stable Floor at age 65 (approximately \$233,000 in 2025 dollars at the general-equilibrium realizable return; Neo-Solon, 2026e, Section 6.7); the deflationary purchasing-power gain accrues to wages and cash rather than to the locked floor.

Mode B — Split-Dividend System — Full-rate 60/40 split; ~0% drift near a one-half transaction ratio. K1 calibrated identically to Mode A. K1 and K2 together calibrated to the full real-growth-matched amount; the post-K1 budget is split sixty/forty between the locked floor (K2) and a standing, spendable K3 dividend. KI inactive. This is the same steady state the system lands in after the debt-retirement transition (Mode T-stable). It builds the largest real Stable Floor of the base Modes — approximately \$413,000 at the general-equilibrium realizable return (Neo-Solon, 2026e, Section 6.7) — alongside the K3 dividend.

Mode C — Inflationary System — Targets approximately 2 percent inflation via price-level path targeting. K1 and K2 operate at Mode A parameters. KI active and annually calibrated to close cumulative CPI gap; distributed monthly to all citizens as unlocked, spendable citizen dividend (~\$108/month at launch at a flat ~1.98% of M2, ~\$115/month within ~2 years, rising with the economy). Produces a smaller real Stable Floor (~\$230K in 2025 dollars) in exchange for current dividend income across working life.

Mode D — Distributed System — Full-rate pure dividend; ~0% drift near a one-half transaction ratio, with no locked floor. K1 inactive; the full real-growth-matched budget is paid entirely as the K3 dividend ($\kappa_d = 100\%$), so the price level holds flat because the dividend reallocates within the issuance budget rather than adding to it. Builds essentially no Stable Floor by design — its benefit is delivered as maximal current dividend income, which a citizen may save or privately invest at the unattenuated ~6.7% return that Mode D leaves intact by deepening no aggregate capital stock (Neo-Solon, 2026e, Section 6.7). The dividend sits exactly on the price-stability leash ($K3 \approx g_r \cdot M^{\wedge T} \approx \230 billion at launch).

Mode Ω — Price-Stability Configuration — The framework’s flagship configuration. A two-sided stabilizing rule that solves for exact price stability in any economy, holding the price level flat in both directions rather than committing to a fixed split (Section 8). Guarantees a stable dollar as a matter of rule rather than discretion; the reference Mode against which the base Modes’ drift is measured.

Mode 0 — Zero-Issuance Configuration — The limiting case in which every creation channel is inactive ($K1 = K2 = KI = 0$) and no new money is created. Recovers the fixed-supply, hard-money corner: with real output growing against a fixed supply, the price level falls at roughly the real-growth rate. Included as evidence that zero is one of the framework’s settings, not a press stuck on “more” (Section 7.5).

Mode Λ — Adaptive Multi-Governor Configuration — An illustrative synthesized Mode combining adaptive K1 multipliers (demographic, population, youth bonus), adaptive K2 boosters (Productivity Governor, Demographic Governor), and a conditional KI that activates only under specified stress conditions. Demonstrates the framework’s extensibility beyond the four base Modes. Not a recommendation; a worked example of the open Mode menu described in Sections 7.4 and 8.

Channel — An issuance mechanism. The framework defines three: K1 (citizenship), K2 (growth), and KI (citizen dividend). Each Channel has its own trigger, calibration formula, destination, and lock rule.

Capital Markets Channel — The transmission mechanism by which K1 and K2 enter circulation. New money funds equity purchases; cash flows to share-sellers; Stable Floor accounts hold equity. Active in all Modes. Produces minimal direct consumer-price inflation pressure.

Direct Deposit Channel — The transmission mechanism by which the circulating-pool channels (K3 and KI) enter circulation: new money deposited directly into citizen transaction accounts as spendable dollars, as opposed to the capital-markets channel that routes K1 and K2 into locked Stable Floors. K3 is active in any Mode with $\kappa_d > 0$; KI is active in Mode C (permanently) and Mode Λ (conditionally).

K1 — Citizenship Endowment — Triggered by each new verified citizen at birth or naturalization. Locked Stable Floor deposit calibrated at 2.5% of GDP per capita. May include adaptive multipliers in extended configurations like Mode Λ .

K2 — Growth Dividend — Triggered by real economic growth. Calibrated annually from the smoothed Composite Productivity Index; distributed monthly to all living citizens' Stable Floor accounts. May include adaptive boosters in extended configurations.

K3 — Consumer Dividend — The growth-indexed citizen-dividend channel. Like KI it deposits unlocked, immediately spendable dollars equally to every citizen, but it is triggered by real economic growth rather than by an inflation gap, and it draws a share κ_d of the same growth budget that funds K2 — routing that share to the circulating pool (current income) instead of the Stable Floor pool (future wealth). Because K1, K2, and K3 partition one growth budget, K3 reallocates issuance between pools without expanding the money supply above the growth-matched line. κ_d is a Mode parameter: zero in the base Modes, positive in a consumer-dividend Mode, bounded above by the circulating-pool price ceiling and by Stable Floor adequacy.

KI — Inflation-Gap Stabilizer — The inflation-gap stabilizer. When the cumulative price level runs below its target path, KI issues new money — deposited equally to all citizens as spendable monthly income — to close the gap. Calibrated annually. Unlocked and immediately spendable. Whether KI produces inflation depends on its calibrated magnitude relative to real economic growth — not on the channel being open. In the base Modes it is permanently active in Mode C (calibrated to produce ~2% inflation) and conditionally active in Mode Λ , but any Mode may activate KI at a lower magnitude without targeting positive inflation.

KI Formula — $KI(t) = \max(0, [\text{target} + \text{real_growth} + \lambda \times \text{gap}(t)] \times M2 - K1_agg - K2_agg)$, where $\text{gap}(t) = \ln(\text{target_price_level}(t-1) / \text{actual_price_level}(t-1))$ and real_growth is total real GDP growth (population growth does not enter the aggregate rate). The maintenance term ($\text{target} + \text{real_growth}$) sustains the target path; the error-correction term $\lambda \times \text{gap}$ closes deviations and vanishes on-path. Closure-speed coefficient λ defaults to 0.5. Self-correcting and non-oscillating.

κ_d (kappa-d) — The consumer-dividend share: the fraction of the real-growth issuance budget paid out through K3 rather than deposited into Stable Floors through K2. A constitutional Mode parameter in $[0, 1)$. $\kappa_d = 0$ recovers the base Modes; a positive κ_d

within the circulating-pool ceiling adds a standing universal dividend while preserving the base Mode's price regime.

Stable Floor — A constitutional minimum capital stake, not a retirement account. Every citizen holds a permanent, individually owned, equity-based position in the national productive economy from the moment of citizenship, funded by citizen seigniorage rather than by taxation or personal savings. The constitutional lock until age 65 is what gives it structural permanence — a floor that can be liquidated at will is not a floor. At 65, citizens may withdraw up to 25 percent as a one-time lump sum and up to 5 percent of remaining balance per year. Heritable to designated beneficiaries under the same withdrawal rules. Constitutionally protected from confiscation. The Stable Floor's significance operates throughout working life, not only at its end: every citizen who knows absolute destitution in old age is constitutionally impossible operates in a fundamentally different economic and psychological environment.

Citizen Equity Stake — Technical name for the Stable Floor. Each citizen owns a transparent, individually titled claim on the productive economy through total-market equity index shares.

Dual Circuit Architecture — Foundational design: a circulating money pool paired with a separately governed Stable Floor pool. Preserved across all Modes.

Composite Productivity Index — Geometric mean of five measures (real GDP per worker, industrial electricity consumption, freight ton-miles, total factor productivity, port and rail throughput) produced by five different agencies. No single agency can manipulate it.

Adaptive Smoothing Rule — Default 5-year rolling average for K2 and KI calibrators, switching to 2-year average at turning points.

Mode Selection — The Tier 2 constitutional choice among available Modes. Requires 67 percent supermajority and 90-day deliberation. Reversible by the same process.

Three-Tier Updatability — Constitutional governance distinguishing Tier 1 (unamendable Model architecture), Tier 2 (parameters and Mode choice amendable by 67 percent supermajority), Tier 3 (operationally updatable details).

FDCA — Federal Digital Currency Authority. Inherits Federal Reserve operational functions. Possesses zero discretionary monetary authority.

Citizen Seigniorage — The value created at the moment new money enters the economy, constitutionally routed to citizens as equity rather than captured by an issuing institution. Classical seigniorage describes profit captured by a sovereign or bank from currency issuance; no established term exists for issuance value that is distributed at creation to the citizens whose economic activity justified it. This paper uses citizen seigniorage to mark that structural difference. Under the current fiat system, this value flows to financial institutions through the Cantillon hierarchy. Under the Citizens Standard, it flows into individually owned citizen equity stakes via K1 and K2, or as spendable income via KI.

Cantillon Effect — The unequal distribution of newly created money where those closest to issuance benefit at full purchasing power before prices adjust. Eliminated internally in all Modes by equal per-citizen distribution at issuance.

Chicago Plan — A 1933 monetary reform proposal for 100 percent reserve banking, validated by the IMF in 2012.

Price-Indexed Debt — Debt contracts denominated in units that adjust annually with the FDCA price index. Required in Mode A; optional in Mode B and Mode C. Modeled on Chile's Unidad de Fomento (1967-present).

Market Exit — Constitutional safeguard providing citizens the right of conversion of Stable Floor balances into non-domestic stores of value under specified conditions.

Surge Brake — Tool 14 in the emergency toolkit. Removes existing money from circulation by retiring it through a graduated, two-tier surcharge — a surge brake plus, under the external interoperability commitment, an anchor-keeping tier — when inflation exceeds the Mode-specific threshold. Because the withdrawn money is retired rather than absorbed into interest-bearing debt, the mechanism accumulates no stock and carries no roll-over interest or $r < g$ sustainability condition.

References

1. Friedman, Milton. *A Program for Monetary Stability*. Fordham University Press, 1960.
2. Benes, Jaromir, and Michael Kumhof. "The Chicago Plan Revisited." IMF Working Paper WP/12/202, August 2012.
3. McLeay, Michael, Amar Radia, and Ryland Thomas. "Money Creation in the Modern Economy." Bank of England Quarterly Bulletin, Q1 2014.
4. Cantillon, Richard. *Essai sur la Nature du Commerce en Général*. London: Royal Economic Society, 1959.
5. Fisher, Irving. "The Debt-Deflation Theory of Great Depressions." *Econometrica*, Vol. 1, No. 4, October 1933.
6. Taylor, John B. "Discretion versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy*, Vol. 39, 1993.
7. Friedman, Milton, and Anna Schwartz. *A Monetary History of the United States, 1867–1960*. Princeton University Press, 1963.
9. Federal Reserve Bank of St. Louis. M2 Money Supply (Seasonally Adjusted). FRED Economic Data Series, March 2026.
11. Department of Commerce. *National Population Estimates*. Census Bureau, Washington, March 2026.
12. Wicksell, Knut. *Interest and Prices (Geldzins und Güterpreise)*. 1898. English translation by R. F. Kahn, Macmillan for the Royal Economic Society, 1936.
13. Fisher, Irving. *Stabilizing the Dollar*. New York: Macmillan, 1920.
14. McCallum, Bennett T. "Robustness Properties of a Rule for Monetary Policy." *Carnegie-Rochester Conference Series on Public Policy*, Vol. 29, 1988.
15. Woodford, Michael. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, 2003.
16. Eggertsson, Gauti B., and Michael Woodford. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, No. 1, 2003.
17. Social Security Administration. *The 2025 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds*. Washington, March 2025.
18. Financial Stability Board. *Global Monitoring Report on Non-Bank Financial Intermediation*. Basel, 2023.
19. Department of Labor. *Consumer Price Index Historical Data*. Bureau of Labor Statistics, Washington, 2025.
20. Department of the Treasury. *Monthly Statement of the Public Debt*. Washington, 2025.
19. Plutarch. *Life of Solon*. In *Lives of the Noble Greeks and Romans*. New York: Modern Library, 2001.

20. Neo-Solon. The Citizens Standard Interactive Engine. Software, 2026. <https://github.com/Neo-Solon/Citizens-Standard>.
21. Neo-Solon (2026b). "The Citizens Standard: A Historical Counterfactual — Empirical Analysis of an Alternative US Monetary Architecture, 1960–2055." Working Paper, June 2026. SSRN 6735078. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6735078.
22. Congressional Budget Office. The Long-Term Budget Outlook. Washington, 2025.
23. Neo-Solon (2026c). "The Citizens Standard: Transition Architecture and Migration Mechanics." Working Paper, June 2026. SSRN 6810741. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6810741.
24. Neo-Solon (2026d). The Citizens Standard: A Statutory Implementation Pathway. Working Paper, 2026. SSRN 6873798. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6873798.
25. Neo-Solon (2026e). The Citizens Standard: A Macroeconomic Model of a Two-Circuit Monetary System. Working Paper, 2026. SSRN 6939418. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6939418.
26. Kydland, Finn E., and Edward C. Prescott. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy*, Vol. 85, No. 3, 1977.
27. Buchanan, James M. "Predictability: The Criterion of Monetary Constitutions." In *In Search of a Monetary Constitution*, ed. Leland Yeager. Harvard University Press, 1962.
28. Brennan, Geoffrey, and James M. Buchanan. *Monopoly in Money and Inflation*. Institute of Economic Affairs, 1981.
29. Hayek, Friedrich A. *Denationalisation of Money*. Institute of Economic Affairs, 1976.
30. Sumner, Scott. *The Money Illusion: Market Monetarism, the Great Recession, and the Future of Monetary Policy*. University of Chicago Press, 2021.
31. Selgin, George. *The Theory of Free Banking*. Rowman & Littlefield, 1988.
32. White, Lawrence H. *Free Banking in Britain*. Cambridge University Press, 1984.
33. Selgin, George, and Lawrence H. White. "How Would the Invisible Hand Handle Money?" *Journal of Economic Literature*, Vol. 32, No. 4, 1994.
34. Goodhart, Charles A. E. *The Evolution of Central Banks*. MIT Press, 1988.
35. Goodhart, Charles A. E. "The Boundary Problem in Financial Regulation." *National Institute Economic Review*, No. 206, 2008.
36. Fisher, Irving. *100% Money*. Adelphi, 1935.

37. Simons, Henry C. "Rules versus Authorities in Monetary Policy." *Journal of Political Economy*, Vol. 44, No. 1, 1936.
38. Friedman, Milton. "A Monetary and Fiscal Framework for Economic Stability." *American Economic Review*, Vol. 38, No. 3, 1948.
39. Tobin, James. "Financial Innovation and Deregulation in Perspective." *Bank of Japan Monetary and Economic Studies*, Vol. 3, No. 2, 1985.
40. Jackson, Andrew, and Ben Dyson. *Modernising Money*. Positive Money, 2012.
41. Sigurjónsson, Frosti. *Monetary Reform: A Better Monetary System for Iceland*. Report commissioned by the Prime Minister of Iceland, 2015.
42. Meade, James E. *Agathotopia: The Economics of Partnership*. Aberdeen University Press, 1989.
43. Bruenig, Matt. *Social Wealth Fund for America*. People's Policy Project, 2018.
44. Hamilton, Darrick, and William Darity. "Can 'Baby Bonds' Eliminate the Racial Wealth Gap?" *Review of Black Political Economy*, Vol. 37, 2010.
45. Blasi, Joseph, Richard Freeman, and Douglas Kruse. *The Citizen's Share*. Yale University Press, 2013.
46. Piketty, Thomas. *Capital and Ideology*. Harvard University Press, 2020.
47. Farmer, Roger E. A. *Prosperity for All: How to Prevent Financial Crises*. Oxford University Press, 2016.
48. Van Parijs, Philippe. *Real Freedom for All*. Oxford University Press, 1995.
49. Mosler, Warren. *Soft Currency Economics*. 1996.
50. Wray, L. Randall. *Modern Money Theory*. Palgrave Macmillan, 2012.
51. Kelton, Stephanie. *The Deficit Myth*. PublicAffairs, 2020.
52. Baker, Malcolm, and Jeffrey Wurgler. "Market Timing and Capital Structure." *Journal of Finance*, Vol. 57, No. 1, February 2002.
53. Jiang, Bin, and Tim Koller. "The Savvy Executive's Guide to Buying Back Shares." McKinsey & Company, October 2011.
54. Bank for International Settlements. *BIS Quarterly Review*, September 2020 (special feature on corporate share buybacks).
55. U.S. Securities and Exchange Commission. *Response to Congress: Negative Net Equity Issuance*. December 2020.

56. Board of Governors of the Federal Reserve System. "Equity Issuance and Retirement by Nonfinancial Corporations." Enhanced Financial Accounts, FEDS Note, June 2017.
57. Goetzmann, William N., and Philippe Jorion. "A Century of Global Stock Markets." NBER Working Paper 5901, 1997 (published as Jorion and Goetzmann, "Global Stock Markets in the Twentieth Century," *Journal of Finance* 54, no. 3 (1999): 953–980).
58. The Vanguard Group. "How America Saves 2025." Valley Forge, PA: Vanguard, 2025.
59. Atkeson, Andrew, and Patrick J. Kehoe. "Deflation and Depression: Is There an Empirical Link?" *American Economic Review* 94, no. 2 (2004): 99–103.
60. Bordo, Michael D., John Landon-Lane, and Angela Redish. "Good versus Bad Deflation: Lessons from the Gold Standard Era." NBER Working Paper 10329, 2004.
61. Lucas, Robert E., Jr. "Inflation and Welfare." *Econometrica* 68, no. 2 (2000): 247–274.
62. Benati, Luca, Robert E. Lucas, Jr., Juan Pablo Nicolini, and Warren Weber. "International Evidence on Long-Run Money Demand." *Journal of Monetary Economics* 117 (2021): 43–63.

Appendix A: Mode Interaction Simulations

This appendix presents three figures that supplement the main paper's analysis. The price-level trajectories in each are computed by compounding that Mode's drift assumption (approximately 1.9 percent deflation in Mode A, price stability in Mode B, and approximately 2 percent inflation in Mode C) at a 7 percent nominal equity return. Figures A.1 and A.2 set their lifetime-value endpoints to the validated full-rate Stable Floor results of the empirical paper (Neo-Solon, 2026b); Figure A.3 computes its lifetime-value comparison directly from the framework's issuance mechanics as implemented in the companion interactive engine. Each citizen tracked is a single representative cohort born at launch. The figures in this appendix demonstrate three architectural properties: that constitutional Mode transitions execute cleanly without breaking the system, that generational Mode rotation produces stable outcomes rather than chaos, and that mid-crisis Mode transitions are mechanically feasible but expensive in lifetime real wealth — making the emergency toolkit (rather than Mode change) the appropriate first-line crisis response.

A.1 Constitutional Mode Transitions

Figure A.1 demonstrates that all six possible Mode transitions at year 30 of a 65-year horizon execute cleanly. The price level reverses trend smoothly when transitioning between Modes. The KI channel activates and deactivates within one year of constitutional ratification. Lifetime real value accumulation differs across scenarios but follows predictable patterns based on which Modes the citizen experienced.

Constitutional Mode Transitions at Year 30 — 65-Year Horizon

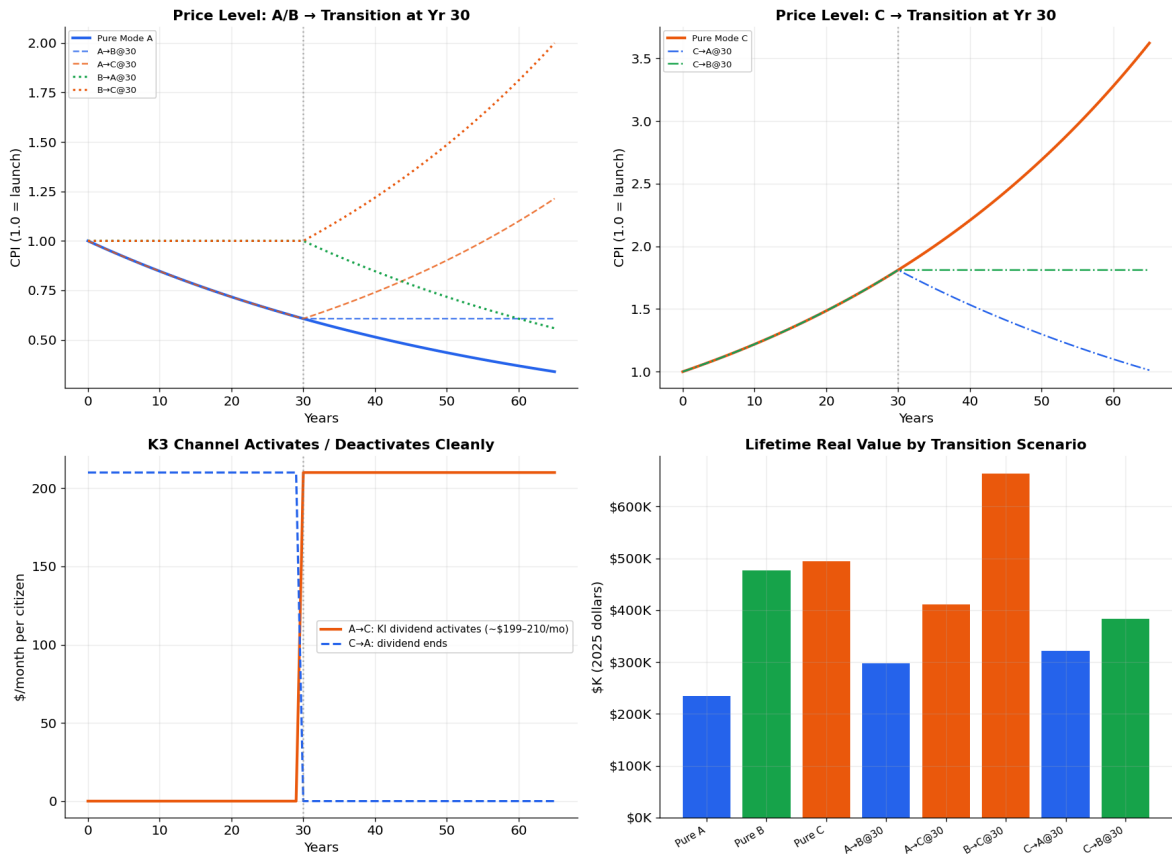


Figure A.1. Constitutional Mode transitions at year 30. All six possible transitions among Modes A, B, and C execute cleanly. Price levels reverse trend smoothly; KI activates and deactivates within one year; lifetime real value follows predictable patterns by Mode sequence.

A.2 Generational Mode Rotation

Figure A.2 simulates a society that ratifies a Mode change every 30 years for 200 years, oscillating between Mode A and Mode C. The framework remains structurally stable across these repeated transitions. The price level traces a sawtooth pattern — declining in Mode A periods, rising in Mode C periods — with each Mode hitting its inflation target within its window. The Stable Floor accumulates monotonically across the entire 200-year horizon despite the repeated regime changes, demonstrating that the architecture’s load-bearing properties survive substantial intergenerational instability in Mode preference.

Generational Mode Rotation — Mode A ↔ Mode C Every 30 Years (200-Year Horizon)



Figure A.2. Generational Mode rotation over 200 years — Modes A and C alternating every 30 years. The price level traces a sawtooth pattern; each Mode hits its inflation target within its window. The framework remains structurally stable across repeated regime changes.

A.3 Mid-Crisis Mode Transition

Figure A.3 examines the consequences of transitioning Mode during economic stress. The simulation imposes a recession (real GDP growth of -2 percent) in years 25-26 and compares two responses: staying in Mode A through the recession versus ratifying a Mode change to Mode C at year 27 (immediately after the recession ends). The Mode-C scenario activates KI and provides immediate citizen dividend support. Because Mode A and Mode C share identical K1 and K2 calibration, the switch does not reduce the locked Stable Floor — both build nearly the same real floor (~\$230,000 for Mode C versus ~\$233,000 for Mode A on the realizable return). What the citizen trades is not floor size but the kind of benefit: Mode A’s mild deflation raises the purchasing power of cash and wages over the horizon, whereas after the switch Mode C delivers value as a visible KI dividend while accepting mild inflation. Measured on total real captured value — the floor plus the cumulative dividend stream Mode C provides after the switch — the move is not a loss but a modest gain, because the dividend adds to what the citizen can hold and spend. The deeper point holds regardless of sign: a Mode change is a permanent regime choice with real, diffuse tradeoffs — the deflationary real-wage benefit given up against the dividend and inflation gained — not a lever to be pulled reactively in a short recession. This supports the main paper’s claim (Section 17.4) that Mode transitions are appropriate as durable constitutional choices while emergency tools are the appropriate response to short-term crises.

Mid-Crisis Mode Transition — Recession at Years 25-26, Mode Change at Year 27

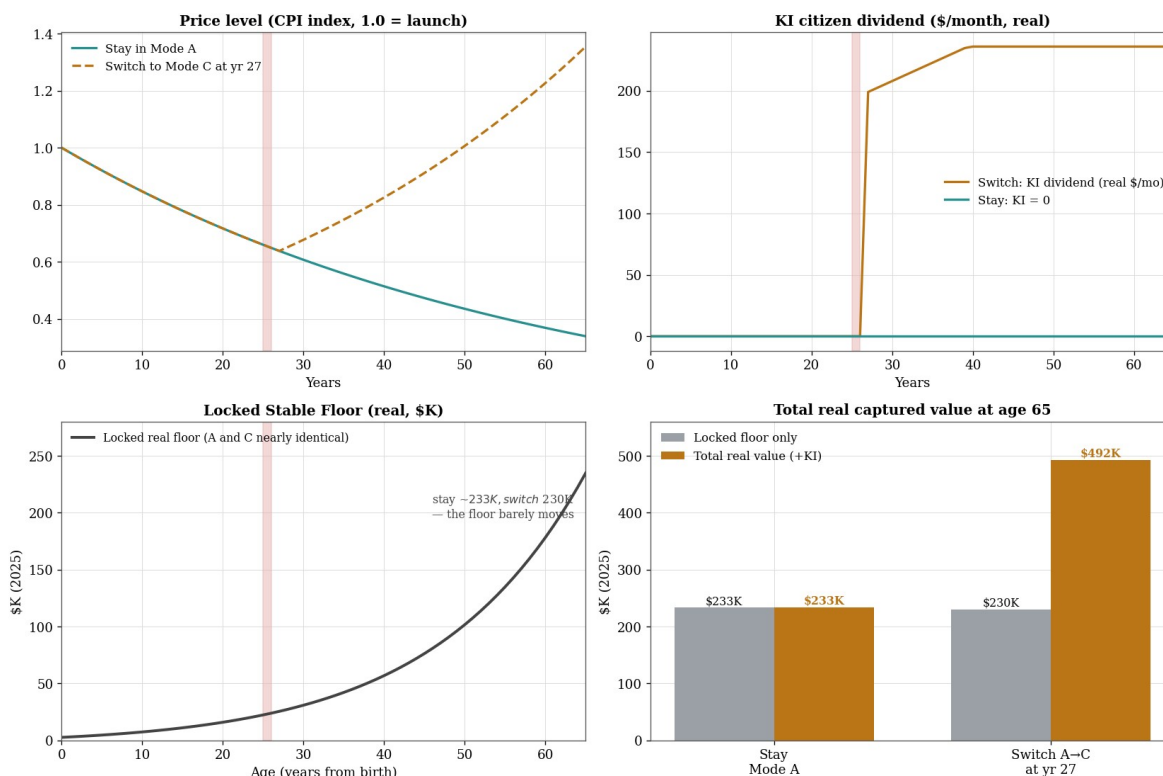


Figure A.3. Mid-crisis Mode transition during a simulated recession (years 25-26), comparing staying in Mode A with switching to Mode C at year 27. Because Modes A and C build nearly the same Stable Floor, the switch leaves the locked floor essentially unchanged (~\$230K); its effect is to add Mode C’s KI dividend, which raises total captured real value to approximately \$310K (the KI stream worth ~\$80K in launch-year real terms, ~\$142K nominal as spent) while giving up Mode A’s deflationary real-wage gains. Emergency tools (Section 12) are the appropriate first-line crisis response; a Mode change is a durable constitutional choice.

A.4 Figure Construction and Provenance

These figures represent illustrative, deterministic baselines rather than stochastic forecasts. They are intended as architectural validation — demonstrating that the framework’s mechanics produce coherent outcomes under reasonable conditions — not as predictions. Real-world implementation would face stochastic variation in equity returns, real growth, inflation, and population dynamics that these figures do not model; the companion empirical paper (Neo-Solon, 2026b) addresses precisely this stochastic dimension under retrospective historical resampling via Monte Carlo. Provenance: Figures A.1 and A.2 are produced by paper1_figures.py and Figure A.3 by paper1_figA3.py in the replication repository (<https://github.com/Neo-Solon/Citizens-Standard>); the latter regenerates the mid-crisis comparison directly from the interactive engine’s issuance mechanics and reports both the locked-floor and total-real-value outcomes. Because Modes A and C build nearly the same real Stable Floor, the locked-floor comparison shows essentially no difference between staying and switching; the

meaningful difference is Mode C's added KI dividend, which only the total-real measure captures.

Appendix B: Interactive Engine

A companion interactive engine accompanies this paper. The engine implements the single-cohort projection model used throughout — K1 issued once at each citizenship event, K2 issued annually as a share of real growth, KI (when active) distributed monthly — and exposes every parameter the framework depends on as an adjustable input. Readers can verify the paper’s numerical claims directly, stress-test the framework under alternative assumptions, and compare the Citizens Standard side by side against seven alternative monetary systems: the Federal Reserve status quo with Social Security, the Chicago Plan, Friedman’s k-percent rule, Bitcoin, the Alaska Permanent Fund Dividend model, tax-funded universal basic income, and Modern Monetary Theory with a federal jobs guarantee.

B.1 Availability

The engine is available in two forms. A standalone web application with sliders, live charts, and a side-by-side comparison view runs in any modern browser without installation. A six-sheet spreadsheet with over a thousand live formulas allows readers to fork the model, modify assumptions, or extend the comparison set. Both implementations share identical underlying mathematics. The engine and its source code are hosted at <https://github.com/Neo-Solon/Citizens-Standard>.

B.2 What the Engine Exposes

The engine surfaces every variable the framework’s outcomes depend on: the issuance channels (K1, K2, K3, KI calibration parameters); the macroeconomic environment (real GDP growth rate, population growth rate, equity index return assumption, CPI inflation rate); the launch conditions (M2 money supply, nominal GDP, population, projection horizon); and for the comparison view, personal savings rate, working-life span, Social Security replacement rate, and UBI/MMT comparison parameters. Four preset Modes (A, B, C, D) load the paper’s reference configurations with one click; a Custom mode is automatically selected as soon as any slider departs from a preset.

B.3 What the Engine Outputs

For a representative cohort projected from year zero to the user-set horizon, the engine produces the Stable Floor balance at horizon in launch-year purchasing power, annual income at the conventional four percent withdrawal rate, the KI monthly dividend per citizen, total annual issuance as a percentage of M2 and GDP, and year-by-year trajectories for M2, GDP, population, the CPI index, and the cohort’s accumulated capital stake. The comparison view shows each system’s locked capital balance at age 65, lifetime cash transfers, annual income at 4 percent withdrawal, funding source, and whether it is rules-based.

B.4 Honest Caveats

The engine carries the same honest caveats as the paper itself. KI is implemented as a flat percentage of M2 rather than the full price-level path targeting formula; the engine reproduces launch-year KI magnitudes but not the dynamic self-correction behavior. Equity returns are treated as nominal; readers who want real returns should set the CPI input to zero. Running the engine at the paper's three preset Mode configurations reproduces approximately sixty-five to eighty percent of the Stable Floor values stated for Modes A and B and approximately one hundred percent for Mode C. The engine is intended for intuition and verification, not for policy calibration.

B.5 Open Release

The engine and its source code are released openly. Readers are encouraged to fork the project, extend the comparator set, modify the underlying assumptions, embed the engine in their own work, or adapt it to other national contexts. Attribution to this paper is appreciated but not required.

The engine exists because the framework's central claim — that monetary outcomes are properly the subject of constitutional politics rather than committee discretion — is best supported by giving readers tools to verify the claim themselves. A paper that stops at tables asks readers to trust the author. A paper that ships a tool asks readers to check the work.

Technical Appendix

The Citizens Standard: One Model, Many Systems — Neo-Solon (2026) · Working Paper · SSRN: 6702518

This appendix provides the complete mathematical specification of the issuance channels (K1, K2, K3, KI), the Stable Floor accumulation model, the cross-Mode calibration and outcome formulas, the KI price-level path targeting formula, and the emergency toolkit ceilings. All parameter values are stated with sources. The companion simulation engine is available at github.com/Neo-Solon/Citizens-Standard.

A.1 Launch Parameters and Shared Variables

All Mode calibrations use the following launch-year values. All projected Stable Floor balances are in December 2025 real dollars unless otherwise stated.

Parameter	Symbol	Value	Source
Nominal US GDP	GDP	\$30,762B	BEA GDPA Apr 2026
US population	N	342M	Census Vintage 2025
GDP per capita	gdppc	\$90,000	GDP / N (paper rounds to \$30.8T / 342M)
M2 money supply	M2	\$22,366B	FRED M2SL end-of-period Dec 2025
Real GDP growth	g_r	2.0%/yr	Base-case assumption (simulation)
Population growth	g_p	0.5%/yr	Base-case assumption
Realizable return on capital	R_real	≈5.4% / 4.3%	GE MPK by Mode (2026e §6.7)
New citizens / year	n	~4.45M	0.4% pop growth × 342M (net additions)
K1 fraction (all Modes)	alpha	2.5%	2.5% of GDP per capita
K2 fraction (Mode A)	beta_A	14.5%	19.5% of M2 growth, less K1
K2 fraction (Mode B)	beta_B	100% of real-growth-matched	Full (g_r × M2), less K1

Parameter	Symbol	Value	Source
K3 consumer share (Mode parameter)	kappa_d	0 in base Modes	Section 3.4; splits the real-growth budget between floors (K2) and dividend (K3)
K1 closure speed (Mode C)	lambda	0.5	Default; closes half of gap per year
K1 inflation target	pi*	2.0%	Constitutionally specified
Accumulation horizon	T	65 years	Birth to retirement age

M2 note. Launch M2 is \$22,366B (exact FRED M2SL end-of-period Dec 2025), which rounds to \$22.4T; this single figure is used throughout, including all emergency-tool ceilings in Section 14. Mode B's full-rate issuance claim in Section 5.1 ($1.0 \times 2\% \times \$22,366B = \sim\$447B$ total issuance; the residual after K1's $\sim\$9B$ is $\sim\$438B$, split 60/40 into K2 floor $\sim\$263B$ and K3 dividend $\sim\$175B$) uses this figure. Total issuance is approximately 2.0% of M2.

Equity-index note. The structural buyer purchases a total-market equity index — operationally the CRSP US Total Market Index (the basis of the largest total-market index funds), or equivalently the Wilshire 5000 Total Market Index, both spanning essentially all US-listed equities. Flow shares are expressed against US total-market equity capitalization, taken as $\sim\$69T$: the Wilshire 5000 Total Market full-cap value, \$69.1T at year-end 2025 (Wilshire / FT Wilshire). This single denominator is used throughout for every “percent of market capitalization” figure; at it the Mode B flow ($\sim\$272B$) is 0.39% and the base-Mode $\kappa_d = 0$ flow ($\sim\$447B$) is 0.65%.

A.2 The Issuance Channels

A.2.1 K1 — Citizenship Endowment (All Modes)

K1 is calibrated identically in Modes A, B, and C:

$$K1(t) = \alpha \times \text{gdppc}(t) = 0.025 \times \text{GDP}(t) / N(t)$$

K1 is a single deposit at birth (or pro-rated at naturalization). It routes to the citizen's locked Stable Floor account.

At launch: $K1 = 0.025 \times \$90,000 = \$2,250$ per new citizen

Annual aggregate: $K1_{\text{agg}} = \$2,250 \times 4.45\text{M new citizens} \approx \9B/yr
(naturalizations pro-rated) = 0.04% of M2

Naturalization pro-ration. A citizen naturalizing at age a receives $(65 - a) / 65$ of the full K1 endowment at the moment citizenship is conferred.

$$K1_{\text{nat}}(a) = K1 \times (65 - a) / 65$$

A.2.2 K2 — Growth Dividend

K2 calibration differs across Modes:

Mode	K2 Formula	Launch value	Total issuance % M2
A	$K2 = 0.175 \times g_r \times M2 - K1$	\$69B/yr (\$203/citizen)	0.31%
B	$K2_agg = 1.0 \times g_r \times M2 - K1_agg$	~\$263B/yr (~\$769/citizen)	~1.95% (K2 only)
C	Same as Mode A	\$69B/yr (\$203/citizen)	0.31% (K1+K2 combined ~0.35%)

Mode A K2 formula derivation. K2 is the residual of the growth-matched envelope (19.5% of M2 real growth) after K1, deposited annually to the Stable Floor:

$$K2(A) = 0.175 \times g_r \times M2 - K1$$

At launch: $0.175 \times 0.02 \times \$22,366B - \$9B = \$69B / 342M = \$203$ per citizen

Mode B K2 formula derivation. Total Mode B issuance is the full real-growth-matched amount. K2 is the residual after K1:

$$Total_B = 1.0 \times g_r \times M2 = 1.0 \times 0.02 \times \$22,366B = \$447B$$

Residual_B = Total_B - K1_agg = \$447B - \$9B = ~\$438B (split 60/40: K2 floor ~\$263B, K3 dividend ~\$175B)

$$K2 \text{ floor}(B)/\text{citizen} = \$263B / 342M = \sim\$769/\text{yr}$$

$$\text{Combined issuance as \% of M2: } \$447B / \$22,366B = 2.0\%$$

Price-stability mechanism. At approximately 2.0% of M2 annual issuance matched to ~2.0% real growth, Mode B holds the price level approximately stable (near 0% annual drift) for economies near a one-half transaction ratio. The money supply grows at the pace of real output, so nominal and real values coincide. This is the same calibration as Mode T-stable.

A.2.3 K3 — Consumer Dividend (All Modes; κ_d -Parameterized)

K3 pays a growth-indexed, equal-per-citizen dividend into the circulating pool. Unlike K1 — which closes an inflation gap and may issue above the growth-matched line — K3 draws from the same real-growth budget that funds K2 and never exceeds it. K1, K2, and K3 partition one budget; the consumer share κ_d sets how much of each year's growth-issuance goes to citizens as current income (K3) rather than into locked Stable Floors (K2):

$$K3_agg = \kappa_d \times (g_r \times M2 - K1_agg)$$

$$K2_agg = (1 - \kappa_d) \times (g_r \times M2 - K1_agg)$$

$K1_agg + K2_agg + K3_agg = g_r \times M2$ (on the growth-matched line)

Worked illustration ($\kappa_d = 0.40$, Mode B launch values). The Mode B growth budget is $g_r \times M2 = \$447B$; $K1_agg \approx \$9B$, leaving a residual of $\$438B$ to split between floors and the dividend:

Residual $R = g_r \times M2 - K1_agg = \$447B - \$9B = \$438B$

$K3_agg = 0.40 \times \$438B = \sim\$175B \rightarrow \$175B / 342M = \sim\$513/yr$
($\sim\$42.75/month$)

$K2_agg = 0.60 \times \$438B = \sim\$263B$ (floors funded at $\sim\$769/citizen/yr$)

A family of four: $\sim\$1,040/yr$ ($\sim\$87/month$)

Why this is price-stable. The launch dividend is modest precisely because K3 is capped by real growth: at $\sim 2\%$ growth the entire budget is $\sim\$447B$, so even two-fifths of it is only $\sim\$175B$ — far below the $\sim\$442B$ that KI issues in Mode C, which is larger exactly because KI is permitted to breach the growth line (that breach is the inflation). K3 instead reallocates within the line, so a positive κ_d shifts dollars from the Stable Floor pool to the circulating pool without expanding the money supply above growth, and the price level is unchanged. The dividend scales over time with M2 and real growth, exactly as K2 does.

Two bounds on κ_d . First, the circulating-pool ceiling: K3 shares the spendable budget with KI and the Stable Floor liquidation flow, so $K3 + KI + liquidation$ must remain within the Y-matched expansion of the circulating pool; a κ_d set beyond it is, by construction, the deliberate step into a higher-inflation Mode rather than a free enlargement of the dividend. Second, Stable Floor adequacy: every point of κ_d diverted from K2 slows floor accumulation, so κ_d is bounded above by the requirement that the projected age-65 floor still clears the constitutional capital guarantee. Within those two bounds, κ_d is a free Mode parameter — zero in the base Modes, positive in a consumer-dividend Mode. The macro model shows the wealth-maximizing choice within those bounds is $\kappa_d = 0$ — the locked floor's diversified compounding dominates a current dividend over a working life — so a positive κ_d does not raise terminal wealth — but it is not therefore valueless: the same model's welfare analysis (Neo-Solon, 2026e, Proposition 8) finds the welfare-optimal κ_d strictly positive under any conventional positive time preference, because a current dividend buys present-consumption utility that locked compounding cannot. A positive κ_d is thus a genuine consumption-smoothing welfare gain traded against terminal wealth, and the $\kappa_d = 0$ default is an explicit accumulation-prioritizing value choice rather than a neutral optimum.

A.2.4 KI — Inflation-Gap Issuance (Mode C; Mode A conditional)

KI uses price-level path targeting to close the cumulative gap between actual and target CPI. The formula as specified in Section 6.1 of the paper:

$$\text{target_price_level}(t) = (1 + \text{pi}^*)^t = (1.02)^t$$

$$\text{gap}(t) = \ln[\text{target_price_level}(t-1) / \text{actual_price_level}(t-1)]$$

$$\text{KI}(t) = \max(0, [\text{pi}^* + \text{g}_r + \text{lambda} \times \text{gap}(t)] \times \text{M2} - \text{K1_agg} - \text{K2_agg})$$

where $\text{lambda} = 0.5$ (default closure speed), $\text{pi}^* = 2\%$ (inflation target), $\text{g}_r =$ real growth, $\text{g}_p =$ population growth, $\text{M2} =$ circulating money supply.

Why ln rather than simple ratio. The natural logarithm handles multi-year compounding correctly: $\ln(\text{desired}/\text{actual})$ produces a consistent closure rate regardless of gap size or duration, whereas a simple percentage ratio overstates the required correction for large accumulated gaps.

Steady state and launch. On the target path $\text{gap}(t) \rightarrow 0$, so the bracket reduces to its maintenance term $\text{pi}^ + \text{g}_r = 2\% + 2\% = 4\%$ of M^T . Net of the K3 dividend (zero in Mode C) and the floor spillover into the transactional circuit ($\sim 0.15\%$ of M^T), KI settles at approximately 3.85% of M^T , about 1.98% of M2 ; the transactional circuit M^T grows 4% against 2% real growth, so under $P = \text{M}^T \cdot V/Y$ realized inflation is exactly $\text{pi}^* = 2\%$, and the price level tracks the target path with no residual offset. On a clean launch the price level begins on its own target path, so $\text{gap} = 0$ and KI issues at this maintenance rate from year one — approximately 3.85% of M^T (1.98% of M2), about \$442B, or roughly \$108 per citizen per month — with no inflation ramp. Worked example (error correction). Suppose instead a year opens with the price level 1% below its target path: then $\text{gap} = \ln(1.01) \approx 0.0100$, the rule adds $\text{lambda} \times \text{gap} \approx 0.5\%$ to that year's issuance, lifting M^T growth to about 4.5% and inflation to about 2.5% until the level is restored; as the gap closes the term decays smoothly to zero and issuance returns to its 4%-of- M^T maintenance level. The overshoot is bounded by lambda — a smaller lambda closes the gap more slowly with less inflation volatility, a larger lambda faster with more. The scholarly lineage of this rule is given in Section 6.1.*

Table A.2. KI ramp to steady state.

Year	KI annual	KI/citizen/month	KI as % M2	Inflation (approx)
Launch (Year 1)	~\$442B	~\$108	~1.98%	~1.0%
Year 3	~\$790B	~\$195	~3.4%	~1.7%
Steady state (Year 5+)	~\$471B (yr ~2)	~\$115	~1.98%	~2.0% (target)

A family of four receives approximately \$432/month at launch, scaling past approximately \$460/month within a couple of years as the dividend tracks a growing transactional circuit at a flat ~1.98% rate.

A.3 Stable Floor Accumulation Model

The Stable Floor SF(i) for a citizen born at launch under each Mode is computed by the companion simulation engine (Appendix B of the paper). This section states the model mechanics and the key parameters that drive the cross-Mode outcomes.

A.3.1 Accumulation Mechanics

Each K1 and K2 deposit is made to the citizen's Stable Floor account, which immediately purchases total-market index shares. Because the account holds equity — a real asset — its value is tracked in real, inflation-adjusted terms, compounding at the real return on the capital it funds, R_{real} , each year. At the universal scale of the program this is not a price-taker market return but the general-equilibrium marginal product of that capital, which attenuates as the citizen stake deepens the stock; it therefore depends on how much of the real-growth line a Mode directs into floors (Neo-Solon, 2026e, Section 6.7), running higher for the low-capture Modes than for the full-budget Mode. It does not depend on the price-level regime the monetary authority targets. The real Stable Floor at retirement is the compounded real value of all deposits:

$$SF_{real}(i) = K1(b) \times (1+R_{real})^T + \sum K2(t) \times (1+R_{real})^{(b+T-t)}$$

where b is birth year, $T = 65$ is the accumulation horizon, and R_{real} is the general-equilibrium realizable return on capital derived in the macroeconomic model (Neo-Solon, 2026e, Section 6.7): approximately 5.4 percent for the low-capture Modes (A, C) and 4.3 percent for the full-budget Mode (B), reflecting the attenuation of the marginal product as the citizen capital share rises. Because deposits and returns are both expressed in real terms, no separate inflation conversion is applied to the floor; the price-level path acts on cash and wages, addressed next.

A.3.2 Price-Level Path and Purchasing Power

The price-level path does not enter the Stable Floor, which is computed in real terms above. It determines the purchasing power of money held as cash and of nominal wages in the circulating pool. The cumulative price level at retirement under each Mode is:

$$\text{Purchasing power of \$1 of cash at Year 65} = 1 / \text{CPI}(T)$$

$\text{CPI}(T)$ is the cumulative price level at retirement under each Mode's inflation path.

Mode	Inflation path	CPI(T) at Year 65	Cash purchasing-power multiplier $1/\text{CPI}(T)$
A	-1.86%/yr deflation	$\text{CPI}(65) = (0.9814)^{65} = 0.294$	3.40×
B	0%/yr drift	$\text{CPI}(65) = (1.000)^{65} = 1.000$	1.00×
C	+2.0%/yr inflation	$\text{CPI}(65) = (1.020)^{65} = 3.623$	0.276×

Why Mode A and Mode C have the same real Stable Floor. K1 and K2 are calibrated identically in Mode A and Mode C — same fractions, same GDP base — so the two accumulate the same real deposits and, capturing the same share of the capital stock and therefore earning the same general-equilibrium return, reach the same real Stable Floor (approximately \$230,000; Neo-Solon, 2026e, Section 6.7). The Modes differ not in the floor but in the circulating pool: Mode A's mild deflation raises the purchasing power of cash

and wages across a working life, while Mode C instead pays the KI dividend. The price-level path acts on that circulating money, not on the locked equity floor.

A.4 Cross-Mode Outcomes — Full Parameter Table

All outcomes for a representative citizen born at launch, under base-case assumptions: R_{real} per Mode ($\approx 5.4\%$ for A/C, 4.3% for B; the general-equilibrium realizable return, Neo-Solon, 2026e, Section 6.7), $g_r = 2\%$, $g_p = 0.5\%$, smooth baseline (no recessions). In 2025 real dollars unless noted.

Metric	Mode A	Mode B	Mode C
Inflation target	-1.86%/yr	0% drift	+2.0%/yr
K1 per new citizen	\$2,250	\$2,250	\$2,250
K2 annual aggregate	\$69B	~\$263B	\$69B
K2 per citizen/year	\$203	~\$769	\$203
KI/K3 per citizen/month	None	~\$42.75 (K3)	\$108 → \$115
Total annual issuance	\$80B (0.35%)	\$447B (2.0%)	\$521B (2.3%)
Stable Floor at 65 — low α/δ band	~\$143K	~\$277K	~\$142K
Stable Floor at 65 — central (GE realizable)	~\$233K	~\$413K	~\$230K
Stable Floor at 65 — high α/δ band	~\$341K	~\$580K	~\$338K
Annual income at 5% real	~\$13.6K/yr	~\$20.7K/yr	~\$13.5K/yr
Cumulative dividend real	—	~\$57K (K3)	~\$80K
Total lifetime real value	~\$233K	~\$470K	~\$310K
Dollar purchasing power $\times 65$	$\times 3.40$ (gain)	$\times 1.00$ (stable)	$\times 0.28$ (-72%)

Method. The Stable Floor holds total-market equity, so its real value is the stream of real K1 and K2 deposits compounded at the real return on capital. At universal scale that return is the general-equilibrium marginal

product of the capital the floors fund, which attenuates as the citizen stake deepens the stock and so depends on each Mode's floor capture (Neo-Solon, 2026e, Section 6.7). Mode A and Mode C share identical K1 and K2 calibration — the same low capture — and therefore earn the same return (approximately 5.4 percent) and build nearly the same real floor, approximately \$233,000 (Mode A) and \$230,000 (Mode C). Mode B directs sixty percent of the real-growth line into the floor; its deeper capture attenuates the return to approximately 4.3 percent, and it builds the largest floor, approximately \$413,000, alongside a standing K3 dividend. The price-level regime acts on cash, wages, and the K1 and K3 dividends, not on the real value of the locked equity floor: Mode A's deflation raises the purchasing power of money held in the circulating pool, not the floor.

A.5 The Composite Productivity Index

The Composite Productivity Index (CPI_prod) is the geometric mean of five independently produced measures. It is the K2 and KI calibrator across all Modes, and is load-bearing for the framework's manipulation-resistance.

Component	Measure	Agency	Cycle
Real GDP per worker	Output per hour, nonfarm business	BEA	Quarterly
Industrial electricity	Total electricity consumption (GWh)	EIA	Monthly
Freight ton-miles	Total freight ton-miles, all modes	BTS	Monthly
Total factor productivity	Multifactor productivity index	BLS	Annual
Port & rail throughput	TEUs + carloads combined index	Census/AAR	Monthly

$$\text{CPI_prod}(t) = [\text{GDP_worker}(t) \times \text{Electricity}(t) \times \text{FreightTM}(t) \times \text{TFP}(t) \times \text{Throughput}(t)]^{1/5}$$

The geometric mean gives each component equal weight in multiplicative terms and is less sensitive to outliers than the arithmetic mean.

Adaptive Smoothing Rule. Default: 5-year rolling average of CPI_prod. Switches to 2-year rolling average when the most recent year's value deviates from the 5-year average by more than 2 percentage points. Reverts to 5-year window once the deviation falls within 1 percentage point. This is fully rules-based — no committee judgment required.

Method A (recommended for K2): use the lower of smoothed GDP growth and smoothed CPI growth, biasing toward conservative issuance when the two measures diverge.

A.6 Emergency Toolkit — Ceiling Summary

All 15 tools have constitutionally specified ceilings expressed as percentages of current M2. This auto-scales capacity as the economy grows and is consistent across all Modes. At launch M2 = \$22.4T (\$22,366B exact; used in all tool ceiling calculations).

#	Tool Name	Ceiling	At launch (\$22.4T)	Failure mode
1	Emergency K1 Provision	2.0%	\$447B	A — Demand collapse
2	Constitutional Rainy-Day Fund	3.5–4.5%	\$795B–\$1.02T	A — Demand collapse
3	Velocity Support Mechanism	—	No new money	Cross-mode
4	Emergency Fiscal Reallocation	1.5–2.2%	\$340B–\$500B	A — Demand collapse
5	Private Credit Guarantee	2.2%	\$500B	C — Banking liquidity
6	Rule-Based Swap Lines	1.3%	\$295B	C — Banking liquidity
7	Immediate Velocity Dividend	0.10%	\$22.4B/yr	A — Demand collapse
8	Deflationary Floor Dividend	0.05%	\$13.4B	A — Demand collapse
9	Enhanced Auto Stabilizers	~0.5%	~\$114B	A — Demand collapse
10	K2 Emergency Boost	Const. bounds	Locked SF only	A — Demand collapse
11	Stable Floor Bridge Loan	10% of bal.	Citizen opt-in	A — Demand collapse
12	Equity Mkt Stability Reserve	1.0%	\$227B	D — Credit cycle
13	Naturalization Processing	Fiscal only	Admin	Cross-mode
14	Surge Brake	3.0%	\$671B	B — Inflation surge
15	M2 Contraction Floor	3.0%	\$671B/yr	A — Demand collapse

#	Tool Name	Ceiling	At launch (\$22.4T)	Failure mode
16	Demurrage Backstop	4.0% cap	Revenue-neutral	A — Deflation hoarding

Total first-year crisis capacity:

Non-inflationary tools excluding Tool 15 (Tools 2, 4, 5, 6, 9, 12, 14): approximately 13–15% of M2 = \$3.0–\$3.5T at launch

Emergency K1 (Tool 1): 2% of M2 = \$447B at launch

M2 Contraction Floor (Tool 15): 3% of M2 = \$671B/yr at launch

Updated total: non-inflationary tools including Tool 15 are approximately 16–18% of M2 (\$3.7–\$4.1T); adding emergency K1 (Tool 1, 2% of M2) gives an approximate grand total of 18–20% of M2

This compares to the genuine liquidity support required in 2008 (~\$2.9T). Tool 15 addresses the Fisher debt-deflation cascade identified in the transition analysis: per the dynamic cascade model in the replication package, the full toolkit offsets approximately 33 percent of a Depression-magnitude M2 contraction and approximately 55 percent of a 2008-magnitude one, with Tool 15 contributing roughly 3 to 4 percentage points in acute scenarios. Partial, not complete — an honest limitation.

Tool 14 trigger calibration by Mode: Mode A: CPI > 5% for 2 consecutive months. Mode B: CPI > 5% for 2 months. Mode C/A: CPI > target + 3pp for 2 months. Tool 15 trigger: M2 contraction > 5% over any rolling 12-month window.

A.7 Full Parameter Table

All baseline values with sources.

Parameter	Value	Source
Population at launch	342M	Census Vintage 2025
GDP per capita at launch	\$90,000	Computed: \$30.8T / 342M
M2 at launch (exact)	\$22,366B	FRED M2SL
M2 at launch (exact, used for tools)	\$22,366B (\$22.4T)	Paper Section 12 convention
K1 fraction (all base Modes)	alpha = 2.5%	Architectural parameter
K1 at launch	\$2,250 per new citizen	0.025 × \$90,000

Parameter	Value	Source
K1 annual aggregate	~\$9B/yr	\$2,250 × 4.45M new citizens
K1 as % M2	0.04%	\$9B / \$22,366B
K2 fraction (Mode A)	beta_A = 19.5% (of M2 growth)	19.5% of M2 growth, less K1
K2 aggregate (Mode A launch)	\$69B/yr	$0.175 \times 2\% \times \$22,366B - \$9B$
K2 per citizen (Mode A)	\$203/yr	\$69B / 342M
Total issuance % M2 (Mode A)	0.35%	$(\$9B + \$69B) / \$22,366B$
Total issuance (Mode B)	\$447B (2.0% of M2)	$1.0 \times 2\% \times \sim\$22.4T \text{ M2}$ (\$22,366B gives \$447B)
K2 aggregate (Mode B)	~\$263B/yr	Total – K1_agg
K2 per citizen (Mode B)	~\$769/yr	\$263B / 342M
K3 consumer share (Mode parameter)	kappa_d = 0 (base Modes)	Section 3.4 / A.2.3; splits the real-growth budget into K2 (floor) and K3
K3 aggregate (Mode B, kappa_d = 0.40)	~\$175B/yr (~\$42.75/citizen/mo)	$0.2 \times (\$447B - \$9B)$; K2 drops to ~\$355B
K1 inflation target	pi* = 2.0%	Constitutionally specified
K1 closure speed	lambda = 0.5	Default; halves gap per year
K1 at launch	~\$442B/yr	Formula at launch conditions
K1/citizen/month at launch	~\$108	$\$442B / 342M / 12$
K1 at steady state	~\$471B/yr	Simulation Year 5+
K1/citizen/month steady	~\$115	$\$471B / 342M / 12$
Realizable return on capital	R_real ≈ 5.4% / 4.3%	GE MPK by Mode (2026e \$6.7)

Parameter	Value	Source
Mode A real SF at 65	~\$233K	Simulation: real return path
Mode A income at 5%	~\$13.6K/yr	5% × \$233K
Mode B real SF at 65	~\$413K	nominal ≈ real (approx. price stable)
Mode B income at 5%	~\$20.7K/yr	5% × \$413K
Mode C real SF at 65	~\$230K	Real return path (same as Mode A)
Mode C income at 5%	~\$13.5K/yr	5% × \$230K
Cumulative KI real (Mode C)	~\$80K	Simulation over 65 years
Total lifetime real (Mode C)	~\$310K	\$230K + \$142K
Tool 1 ceiling (Emergency K1)	2.0% of M2 = \$447B	Section 14.1
Tool 2 ceiling (Rainy Day)	3.5–4.5% of M2	Section 14.1
Tool 12 ceiling (Equity Res.)	1.0% of M2 = \$227B	Section 14.4
Tool 14 ceiling (Sterilize)	3.0% of M2 = \$671B	Section 14.2
Total non-inflationary cap.	~13–15% of M2 = \$3.0–3.5T	Section 14.7

A.8 Mode A Governor Derivation

Mode A's Section 9.6 figures are deterministic outputs of the same Stable Floor accumulation model used for Modes A–D (Appendix A.3), with the issuance channels set by Mode A's constitutional governors rather than by fixed targets. This section states the parameterization so the scenario floors in Table 8 are reproducible; the full implementation is provided as `mode_omega.py` in the architecture replication package.

The base configuration captures 60 percent of the real-growth-matched budget ($K2 = 0.60$, between Mode A's 0.175 and Mode B's 1.0), with K1 at the standard 2.5 percent of GDP per capita. Inflation is derived from issuance exactly as for the base Modes — at 60 percent capture the implied drift is approximately -0.8 percent, a mild deflationary bias — and combined issuance is held below the 3.5-percent-of-M2 ceiling. With

governors inactive, this base produces a real Stable Floor at age 65 of approximately \$403,000. The return is no longer a free assumption: at 60 percent capture the general-equilibrium return on the deepened citizen stock settles near 4.1 percent (Neo-Solon, 2026e, Section 6.7), so the floor sits close to Mode B's, since both run near the same capture.

Under negative population growth the demographic governor engages. Two effects combine: the falling citizen count raises each citizen's share of the K2 budget (a +20-percentage-point lift at -0.5 percent population growth), and the governor itself raises K1 to 1.3× base (the +0.3× setting at the bottom of the constitutional +0.3× to +0.8× range) and the K2 capture by +30 percent (within the +15 to +40 percent range). Together these lift capture to roughly 78 percent; but deeper capture attenuates the general-equilibrium return — from about 4.1 to 3.4 percent — so the net effect is a +19 percent lift rather than the larger figure a fixed-return projection would imply, raising the central floor to approximately \$480,000. The same governor path carries across the three negative-population scenarios, which differ by where the marginal product lands within the α/δ robustness band. In the productivity-boom scenario the productivity governor adds a further +30 percent to K2 — its constitutional ceiling — stacking with the demographic governor and deepening capture to about 96 percent.

Applying these settings yields the five Table 8 floors — approximately \$403,000 (normal), \$480,000 (negative pop, central general-equilibrium return), \$638,000 (negative pop, high α/δ band), \$343,000 (negative pop, low α/δ band), and \$541,000 (productivity boom) — with the 5-percent withdrawal incomes shown alongside. Two features are worth noting. At normal demographics Mode Λ lands alongside Mode B (\approx \$403,000 versus \approx \$413,000), since both run near 60 percent capture. And at the low end of the α/δ band it edges above Mode B (\approx \$343,000 versus \approx \$277,000), because the demographic governor outweighs Mode B's fixed capture when the marginal product is low. The broader lesson is that the governor and the return work against each other: lifting capture to defend the floor against demographic decline also attenuates the realizable return, so Mode Λ 's outcomes compress into a tighter band than a fixed-return projection would suggest. Every multiplier above is constitutionally fixed and formula-triggered; the calibration choices are recorded in `mode_omega.py` and can be varied within the published governor ranges.