

THE CITIZENS STANDARD: TRANSITION ARCHITECTURE AND MIGRATION MECHANICS

Neo-Solon

Neo-Solon@hotmail.com

WORKING PAPER — 2026

COMPANION PAPERS

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

The Citizens Standard: A Historical Counterfactual (Neo-Solon, 2026b) · SSRN 6735078

The Citizens Standard: A 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

Abstract

The Citizens Standard (Neo-Solon, 2026a) proposes a constitutional monetary architecture whose destination is well-specified. This paper addresses what the architectural paper defers: the path from the current discretionary monetary system to that destination. The central finding is that the Citizens Standard does not require a monetary revolution to begin. Its most powerful mechanism — universal locked equity compounding — can be launched today as a parallel sovereign wealth layer within the existing monetary system, requiring no Federal Reserve replacement, no constitutional amendment, and no banking restructuring at inception.

We specify a five-phase migration architecture in which each phase is self-contained, generates observable evidence, and creates the institutional conditions for the next phase. The phases span approximately 40 to 60 years from launch to full constitutionalization. We then address quantitatively the four hard transition problems the architectural paper acknowledges but does not resolve: existing government debt conversion mechanics, banking separation and credit stability under phased reserve requirements, equity valuation effects of universal systematic ownership flows, and the timing of constitutional lock credibility.

The quantitative foundation draws on the empirical paper's (Neo-Solon, 2026b) decomposition of the Stable Floor balance: deposited principal — the cumulative K1 and K2 issuance — accounts for only about 5 percent of the final balance, with the remaining roughly 95 percent generated by equity compounding over the locked, fee-minimized, full-horizon accumulation window. The transition implication is that the binding determinant of long-run outcomes is the compounding architecture rather than the scale of monetary issuance in any single year: Phase 1 can launch at near-zero monetary scale while establishing the structure that produces the framework's long-run outcomes, with issuance scaled up over later phases. Full Mode B parameters — K2 calibrated at the full real-growth-matched rate (approximately stable prices near the one-half transaction-ratio balance point), with M2 growing at approximately 2.5 percent annually — are required to deliver the Stable Floor documented in the empirical paper, which against historical US data reaches approximately \$210,000 for the earliest cohort on the realizable basis.

A central quantitative finding concerns the debt transition. The relevant debt is the \$31.4 trillion held by the public (102 percent of GDP), not the \$39 trillion gross total — the \$7.6 trillion of intragovernmental debt is non-marketable and nets out. At enactment this public debt transfers to a Legacy Debt Trust, a wound-down vehicle that may refinance but never expand the stock, resolving the annual rollover wall (roughly one-third of the stock matures each year) without permitting new government borrowing. The transition runs under Mode T: citizen K1 and K2 flow uninterrupted at the full real-growth-matched rate (approximately stable prices near the one-half transaction-ratio balance point), while a transition-only channel, KT, issues money calibrated to a price-level path and directs it to bond redemption rather than to citizens. Because redemption is an asset swap absorbed by a reinvesting holder base, KT retires debt while remaining consumer-price neutral; it is self-throttling on inflation and self-extinguishing once debt stabilizes. Under this path, public debt-to-GDP falls from 102 percent at enactment to approximately 84 percent by Year 10 and 58 percent by Year 20, stabilizing within a moderate operational band of roughly 30 to 60 percent of GDP (central path approximately 45 percent, reached by Year 26) — retiring the debt as a fiscal burden while retaining a standing stock of sovereign bills sized to supply the financial system's safe-asset benchmark and the base for symmetric KT operations — against the CBO's March 2025 projection of 156 percent by 2055 under

current law. The band is the welfare-optimal endpoint identified by a stochastic debt-sustainability analysis calibrated to the interest-growth-differential literature (Blanchard 2019; Mauro and Zhou 2020; Lian, Presbitero and Wiriadinata 2020): because standing debt is near self-financing while the safe rate sits below the growth rate, retiring below the band would spend citizen seigniorage to no sustainability purpose, while the band stays well clear of the debt levels at which crisis risk rises materially. Within the band, KT routes the growth-matched seigniorage to citizen Stable Floors by default and to redemption only as needed to hold the band. Once the band is reached the system lands in Mode T-stable, a permanent steady state that is Mode Ω , holding the price level stable for any transaction ratio (coinciding with the fixed Mode B split at the one-half ratio, and given by the Ω solve otherwise), in which KT remains available as a symmetric open-market instrument for two-sided monetary control.

Stress tests against three migration failure scenarios — political reversal, equity market disruption, and credit contraction — show that the phase architecture contains each failure mode without requiring full restart. International analogues from Chile, Norway, Estonia, and Singapore provide empirical grounding for each phase's feasibility.

JEL classification: E42, E51, E52, E58, E63, H63

Keywords: Citizens Standard; monetary transition; debt retirement; macroeconomic stabilization; reform sequencing; institutional migration

Contents

1. Introduction.....	5
2. The Migration Principle.....	5
2.1 The Sovereign Wealth Precedent.....	6
2.2 What the Migration Principle Implies for Phase Sequencing.....	6
2.3 The Two-Ballot Ratification Sequence.....	6
3. The Five-Phase Migration Architecture.....	7
3.1 Phase 1 — Sovereign Wealth Foundation (Years 1–15).....	7
3.2 Phase 2 — Voluntary Expansion and K2 Scaling (Years 10–25).....	8
3.3 Phase 3 — Banking Separation (Years 20–40).....	8
3.4 Phase 4 — Rule-Based Monetary Layer (Years 35–50).....	8
3.5 Phase 5 — Constitutionalization (Years 45–60).....	9
4. The Four Hard Transition Problems.....	9
4.1 Existing Government Debt Conversion.....	9
4.2 Banking Separation and Credit Stability.....	13
4.3 Equity Valuation Effects of Universal Ownership Flows.....	14
4.4 Constitutional Lock Timing.....	14
5. Failure Mode Analysis.....	15
6. Political Economy of the Migration Path.....	15
7. International Transition Analogues.....	16
8. Discussion.....	17
8.1 Social Security and the Citizens Standard: Consolidation as a Phase 4–5 Question.....	17
9. Conclusion.....	19
References.....	21
Technical Appendix.....	23
A.1 Notation and Shared Variables.....	23
A.2 Sovereign Debt Transition Model (Section 4.1).....	24
A.3 Banking Separation and Credit Stability Model (Section 4.2).....	26
A.4 Equity Valuation Flow Model (Section 4.3).....	27
A.5 Constitutional Lock Credibility Model (Section 4.4).....	29
A.6 Full Parameter Table.....	30

1. Introduction

Monetary reform proposals are evaluated on two distinct questions that are rarely separated cleanly. The first is whether the destination architecture is internally coherent — whether the proposed system, once operational, would function as described. The second is whether a credible path exists from the current system to that destination. These questions have different methods, different literatures, and different failure modes. A framework can be theoretically coherent and transitionally impossible. It can also be transitionally straightforward while theoretically contested. Treating them as the same question obscures more than it clarifies.

The Citizens Standard (Neo-Solon, 2026a) is primarily an answer to the first question. The architectural paper specifies a dual-circuit monetary architecture, a citizen-anchored issuance mechanism, a separated banking structure, and a constitutional governance framework. The empirical paper (Neo-Solon, 2026b) evaluates the framework's retirement security claims against US historical data from 1960 to 2025 and finds that Mode B reliably produces a Stable Floor approximating median actual retirement outcomes across four cohorts on the realizable basis. What neither paper addresses in full is the second question: given that the destination is specified and its retirement-security claims are empirically grounded, how does a society actually get there from a sovereign debt load of roughly \$39 trillion gross, of which approximately \$31.4 trillion is held by the public, a \$22.4 trillion M2 supply created almost entirely through bank credit, a Federal Reserve with an \$6.7 trillion balance sheet, and institutional actors with strong incentives to preserve the current architecture?

This paper addresses that question directly. The central finding can be stated simply: the Citizens Standard does not require a monetary revolution to begin. The framework's most powerful mechanism — the Stable Floor — is an equity ownership structure, not a monetary instrument. A Phase 1 Citizens Standard operating at minimal monetary scale establishes the compounding architecture that produces long-run outcomes without requiring any change to the Federal Reserve, the Treasury market, the banking system, or the constitutional structure. It resembles, from the outside, a sovereign wealth program. Politically, that is its entry point.

The migration from that entry point to the full constitutional architecture is a phased process spanning approximately 40 to 60 years. Each phase is self-contained: it has defined entry conditions, defined success criteria, and defined exit conditions. No phase requires the next phase to succeed. The phases are additive improvements, not steps in an all-or-nothing sequence.

Section 2 develops the migration principle and its sovereign wealth analogues. Section 3 specifies the five-phase architecture. Section 4 addresses the four hard transition problems quantitatively. Section 5 stress-tests the phase architecture. Section 6 examines the political economy. Section 7 draws on international analogues. Section 8 discusses what the analysis establishes and what it does not. Section 9 concludes.

2. The Migration Principle

The Citizens Standard's transition problem is unusual among monetary reform proposals because its most valuable component is not monetary. The Stable Floor — universal,

locked, equity-based ownership funded by seigniorage — is an ownership structure that equity markets can host, that existing custodial infrastructure can administer, and that statutory law can authorize without constitutional amendment. The monetary reform elements of the framework are architecturally necessary for the full system but are not prerequisites for the Stable Floor to begin operating and compounding.

This separability is the migration principle. The transition question is not 'how do we replace the monetary system' but 'how do we begin a sovereign wealth migration that the current system can host, and under what conditions does that migration eventually become the system.'

2.1 The Sovereign Wealth Precedent

Universal equity ownership funded by resource or seigniorage flows is not a novel institutional form. Three existing programs provide direct empirical grounding for Phase 1 feasibility.

The **Alaska Permanent Fund** was established by constitutional amendment approved by Alaska voters in 1976, with the Alaska Permanent Fund Corporation created in 1980 and the first dividend distributed in 1982. As of June 30, 2024, the fund held \$80.5 billion in assets — approximately \$110,000 per Alaska resident. Its governance structure, constitutional protection, and universal distribution are direct precedents for the Citizens Standard's K1 and K2 mechanisms.

The **Norwegian Government Pension Fund Global**, established in 1990, now holds approximately \$2.0 trillion in assets — roughly \$390,000 per Norwegian citizen. Norway's fund demonstrates that sovereign equity ownership at national scale is administratively manageable, that long-duration equity investment produces the compounding outcomes the Citizens Standard projects, and that constitutional governance of a sovereign wealth structure is politically durable across governments and decades.

Singapore's Central Provident Fund, established in 1955 and retained after independence in 1965, mandates universal participation, locks contributions until retirement age, and as of December 2024 managed SGD \$609.5 billion for 4.2 million account holders. The CPF demonstrates that mandatory locked accumulation is politically sustainable and that universal participation produces the behavioral advantages the empirical paper (Neo-Solon, 2026b) identifies as the source of the framework's median-beating record.

2.2 What the Migration Principle Implies for Phase Sequencing

The migration principle has three direct implications. First, Phase 1 requires no Mode vote — the sovereign wealth layer can be authorized by ordinary statute. The Mode selection question, requiring a 67 percent supermajority Tier 2 constitutional ratification, is deferred until citizens are voting on a system with observable performance data. Second, each phase's success criteria are defined in terms of what that phase alone accomplishes. Third, the migration is reversible at every phase boundary except constitutionalization — a reform that can be stopped is easier to start than one that cannot.

2.3 The Two-Ballot Ratification Sequence

Ballot 1 occurs at Phase 1 launch. Citizens authorize a federal sovereign wealth program through ordinary statutory legislation — universal equity accounts for all newborn citizens,

funded by a small seigniorage allocation, invested in a total-market index, locked until age 65. The Federal Reserve continues operating without modification.

Ballot 2 occurs at the Phase 2 to Phase 3 transition, after approximately 10 to 15 years of Phase 1 operation. Citizens vote on whether to constitutionalize and expand what they have already watched work — ratifying the Model, selecting a Mode, and beginning the banking and monetary architecture transition. This is the Tier 2 supermajority vote the architectural paper specifies, and it is more legitimate than a Day 1 constitutional vote because it is informed by evidence rather than projection.

This sequence has a historical analogue. The Social Security Act was enacted in 1935 as ordinary legislation; regular monthly benefits did not begin until 1940. Medicare followed in 1965 after three decades of demonstrated Social Security operation. The Citizens Standard's two-ballot sequence follows this precedent: establish the program through statute, demonstrate its value, then seek the constitutional protection that makes it permanent.

3. The Five-Phase Migration Architecture

The migration proceeds through five phases, each with defined entry conditions, a primary objective, success criteria, and exit conditions. The phases overlap — Phase 2 begins before Phase 1 exits — and each is self-contained. The quantitative parameters use Mode T as the transition configuration throughout, with full-rate K2 funding citizen Stable Floors and the KT channel retiring legacy debt; the system lands in Mode T-stable once the debt stabilizes within its operational band.

Table 1. Five-phase migration timeline. Years are approximate from Phase 1 launch. Phase overlaps are deliberate. K2 level is percentage of full Mode B calibration.

Phase	Years	Primary Objective	K2 Level	Key Milestone
1	1–15	Universal K1 + minimal K2	0–10%	First cohort compounding
2	10–25	K2 scaling + voluntary expansion	10–50%	Constitutional ratification
3	20–40	Banking separation + full K2	50–100%	Full Mode T operational
4	35–50	Fed conversion + rules-based layer	100%	FDCA sole monetary authority
5	45–60	Constitutional lock	100%	Tier 1 provisions ratified

3.1 Phase 1 — Sovereign Wealth Foundation (Years 1–15)

Entry conditions: Simple majority legislation. No constitutional amendment, no Federal Reserve modification, no banking restructuring.

Primary objective: Establish universal equity ownership for all newborn citizens. Every child born in the United States from Phase 1 launch receives a K1 deposit — approximately \$2,250 at launch — invested automatically in a total-market index fund and

locked until age 65. A small K2 stream at 5 to 10 percent of the full-rate K2 is introduced in Year 3. Annual K2 at 10 percent is approximately \$45 billion — 0.15 percent of GDP.

Quantitative outcomes: A citizen born at Phase 1 launch with K1 plus 10 percent K2 accumulates approximately \$100,000 in 2025 real dollars at age 65 on the realizable basis — below the \$260,000 median actual retirement benchmark from the empirical paper (Neo-Solon, 2026b). Phase 1 establishes the compounding trajectory that subsequent phases build on.

Success criteria: Universal enrollment operational; first cohort compounding trajectory within 10 percent of model projections; no material inflation effect; political durability through at least one election cycle.

3.2 Phase 2 — Voluntary Expansion and K2 Scaling (Years 10–25)

Primary objective: Scale K2 from 10 to 50 percent of full Mode B, with annual issuance rising from approximately \$45 billion to \$219 billion. Voluntary opt-in mechanisms allow pre-Phase-1 citizens to establish Stable Floor accounts through tax-advantaged contributions and partial Social Security integration.

Quantitative outcomes: A citizen born at Phase 2 midpoint accumulates approximately \$175,000 in 2025 real dollars at age 65 — approaching but not yet exceeding the \$260,000 median actual benchmark.

Success criteria: K2 at 50 percent sustained for 3 years without material inflation effect; voluntary opt-in above 20 percent of eligible cohorts; Ballot 2 constitutional ratification completed; banking separation legislation introduced.

3.3 Phase 3 — Banking Separation (Years 20–40)

Primary objective: Implement the payment-credit separation specified in the architectural paper. Transaction accounts migrate to full-reserve status over 20 years in three tranches. The Transition Lending Facility operates throughout. K2 scales from 50 to 100 percent of the full rate during Phase 3.

Quantitative outcomes: A citizen born at Phase 3 midpoint accumulates approximately \$330,000 in 2025 real dollars at age 65 — above the median actual benchmark. The full Mode B Stable Floor documented in the empirical paper (approximately \$210,000 for the earliest historical cohort on the realizable basis, rising toward approximately \$570,000 for forward cohorts accumulating against a larger economic base) is achieved for citizens benefiting from the complete 65-year compounding window.

Success criteria: Full-reserve transaction accounts at all major institutions; reserve requirements at 100 percent; TLF wound down; K2 at 100 percent of full Mode B for 2 years; no recession attributable to credit contraction.

3.4 Phase 4 — Rule-Based Monetary Layer (Years 35–50)

Primary objective: Transfer Federal Reserve discretionary authority to the FDCA's constitutional protocol. Payment system functions transfer to the FDCA, open market operations cease, and the supervisory role is constitutionally bounded. The fifteen-tool emergency toolkit replaces lender-of-last-resort discretion.

3.5 Phase 5 — Constitutionalization (Years 45–60)

Primary objective: Achieve the Tier 1 constitutional lock. Four decades of performance data exist, the oldest Phase 1 cohorts are approaching retirement, and constitutional ratification confirms what is already operationally demonstrated. The lock removes the Stable Floor from ordinary legislative bargaining permanently. The Mode ratified at this stage is the full calibration vector — the K1, K2, K3, and KI settings together with the consumer-dividend share κ_d (architectural paper, Neo-Solon 2026a, Section 3.4) — adopted as one package; once Mode T sunsets, a Mode with $\kappa_d > 0$ brings the permanent consumer-dividend channel K3 online alongside the floor-building channels, while $\kappa_d = 0$ lands the system in a pure floor-building steady state.

4. The Four Hard Transition Problems

4.1 Existing Government Debt Conversion

The United States carried approximately \$39.0 trillion in gross federal debt as of early 2026, with annual net interest approaching \$1 trillion. The relevant figure for the transition, however, is not the gross total but the debt held by the public — approximately \$31.4 trillion, or 102 percent of GDP. The remaining \$7.6 trillion is intragovernmental debt: special-issue, non-marketable securities the government owes its own trust funds (Social Security OASI, federal and military retirement, Medicare Hospital Insurance). As the Committee for a Responsible Federal Budget states, intragovernmental debt “has no net effect on the government’s overall finances.” It does not trade in the Treasury market, carries no rollover risk, and is redeemed internally as the programs draw down. The Legacy Debt Trust therefore needs to refinance only the \$31.4 trillion held by the public; the trust-fund obligations are addressed separately by the Social Security consolidation question discussed in Section 8.1.

The rollover wall. The genuine challenge is not the headline ratio but the refinancing schedule. With an average maturity near six years, approximately one-third of the public debt — on the order of \$10 to 12 trillion — matures each year and must be refinanced. Against federal revenues near \$5 trillion, this is the constraint that makes a simple “no new debt” rule impossible to impose overnight. The transition must refinance the maturing stock without either defaulting or forcing an abrupt fiscal contraction.

The Legacy Debt Trust. At enactment, all pre-existing public debt transfers to a legally distinct vehicle, the Legacy Debt Trust. The Trust is the sole entity permitted to refinance maturing pre-enactment securities; it may issue new instruments to replace maturing ones but may never expand the stock. This resolves the rollover wall without violating the no-new-debt constraint on the government itself: the government issues nothing new, while the Trust winds the legacy stock down. As each tranche matures and is refinanced, its coupon falls from the legacy average of approximately 4.5 percent toward the post-transition yield of approximately 3.0 percent nominal (about 1.5 percent real under price stability), so that by Year 6 the average coupon has largely repriced and the interest-growth spread has collapsed.

Mode T and the KT channel. The transition operates under Mode T (architectural paper, Section 8A). K1 and K2 fund citizen Stable Floors uninterrupted, with K2 at the full real-growth-matched rate so that consumer prices are held stable. The transition-only KT channel issues money — calibrated to a price-level path at approximately 1.5 percent of

M2 — directed to the Legacy Trust for bond redemption rather than to citizens. Because redemption is an asset swap (the bondholder exchanges a bond for cash and, given a holder base dominated by foreign central banks, pension funds, and the central bank, predominantly reinvests), KT retires debt without the consumer-price shock that distributing the same magnitude to households would produce. The channel is self-throttling: calibration to a price-level path means issuance automatically reduces if consumer inflation rises.

Trajectory. Under Mode T with a primary surplus phasing in from zero to approximately 1.5 percent of GDP over twenty-five years, and KT operating at 1.5 percent of M2, public debt-to-GDP falls from 102 percent at enactment to approximately 84 percent by Year 10 and 58 percent by Year 20, then stabilizes within a moderate operational band of approximately 30 to 60 percent of GDP — central path approximately 45 percent, reached by Year 26 — where it holds rather than continuing to descend (Figure 1). The band is the welfare-optimal endpoint identified by the debt-sustainability analysis of Section 4.1.1. This compares with the Congressional Budget Office's March 2025 projection of 156 percent by 2055 under current law. The KT channel accounts for approximately \$12 trillion of cumulative redemption over the period; the remainder of the descent is delivered by the primary surplus and by nominal output growth. Within the band the standing stock is a deliberate, seigniorage-funded holding — supplying the safe-asset benchmark and the base for symmetric KT operations, not a residual burden — and as the debt enters the band KT throttles down, so the growth-matched seigniorage that is no longer needed to hold position is routed to citizen Stable Floors. Citizen K1 and K2 flows are untouched throughout.

Self-extinguishing and the soft landing. The KT channel retires the public debt — the \$31.4 trillion held outside the government — not to literal zero but down into a moderate operational band of approximately 30 to 60 percent of GDP, with a central path of approximately 45 percent reached by Year 26, where the stock then holds. The band is the welfare-optimal endpoint established by the debt-sustainability analysis of Section 4.1.1. What KT eliminates is debt as a fiscal burden: the rollover wall, the roughly \$1 trillion in annual net interest, and the structural dependence on continuous new borrowing. What it retains is a standing stock of short-term sovereign bills, held by choice rather than necessity — sized for monetary operations and safe-asset supply, not for deficit financing, and fully covered by the same seigniorage rule that funds citizen Stable Floors. This retained stock is not idle. A modern financial system requires a risk-free benchmark asset: money-market funds, repo and collateral chains, and pension and insurance reserves all price off sovereign paper, and a literal-zero sovereign stock would remove that anchor — a problem the United States itself anticipated in the late-1990s surplus projections, when the prospective disappearance of the Treasury market was treated as a cost rather than a triumph. The moderate band preserves the benchmark with room to spare, and held mostly as non-interest-bearing bills it supplies that benchmark without imposing a standing transfer to bondholders. KT does not attempt to retire the intragovernmental debt, which it never held: the roughly \$7.6 trillion of intragovernmental balances (approximately 8 percent of grown GDP by the time the band is reached) is addressed separately by the Social Security consolidation of Section 8.1. Once the band is reached the configuration becomes Mode T-stable — but KT is not so much retired as reoriented. It remains available as a symmetric, rules-based open-market instrument. In normal conditions it is dormant, the stock self-maintaining within the band as a share of GDP, with the growth-matched seigniorage routed to citizen Stable Floors; when consumer inflation runs above target, KT reverses — issuing bills above the band to drain transactional money, then

unwinding once inflation subsides. This reverse-KT drain is the contractionary, QT-equivalent tool that gives the steady state two-sided monetary control, and it resolves an asymmetry in the baseline design, in which the citizen issuance channels expand the money supply but no symmetric instrument contracts it. Because it operates only on transactional balances and never on the locked Stable Floors, price control is achieved without touching citizen capital. The conditional KI_T damper, by contrast, is withdrawn once banking separation completes. None of this is visible to citizen Stable Floors, since KT never deposited into them; and the society may, through the standard Tier 2 supermajority process and informed by decades of observed performance, adopt a different steady-state Mode or remain in Mode T-stable permanently. Mode selection is thus never a debt-management constraint.

Asset-price honesty. Mode T holds consumer prices stable (KT is a price-neutral asset swap, and the citizen-facing split is set for zero derived inflation, coinciding with the fixed Mode B split at the one-half transaction ratio and given by the Mode Ω solve otherwise) but is not neutral with respect to asset prices. Bondholders receiving redemption reinvest, and as sovereign instruments grow scarcer a portion of that capital — estimated at 15 to 35 percent, with a central estimate near 17 percent once the large reinvesting holders are accounted for — rotates into equities and real estate, exerting mild upward pressure on valuations. Combined with Stable Floor purchases, the additional equity demand peaks at approximately 0.75 percent of market capitalization annually, well within the range of existing systematic flows such as retirement-account contributions. Because citizen Stable Floors hold broad-market equity, this appreciation accrues substantially to citizens. The effect is treated quantitatively, including its modest intergenerational return-compression consequence, in Section 4.3.

Table 2. Public debt-to-GDP trajectory under Mode T: Legacy Debt Trust refinancing \$31.4 trillion of public debt, with the KT channel directing approximately 1.5 percent of M2 annually to bond redemption and a primary surplus phasing in to 1.5 percent of GDP over twenty-five years. Public debt is retired down into a moderate operational band of approximately 30 to 60 percent of GDP (central path ~45 percent, reached by Year 26), where it stabilizes rather than falling toward zero. Coupon reprices from the 4.5 percent legacy average toward the ~3.0 percent post-transition yield by Year 6.

Year	Mode	D/GDP	Avg Coupon	Notes
0	T	102%	4.5%	Launch — public debt to Legacy Trust
10	T	84%	3.0%	KT retiring public debt
20	T	58%	3.0%	KT + surplus retiring debt
30	T	45%	3.0%	Band centre reached (~Yr 26); KT throttles
40	T	45%	3.0%	Holding within band; KT dormant
45	T	45%	3.0%	Mode T-stable; seigniorage to citizens

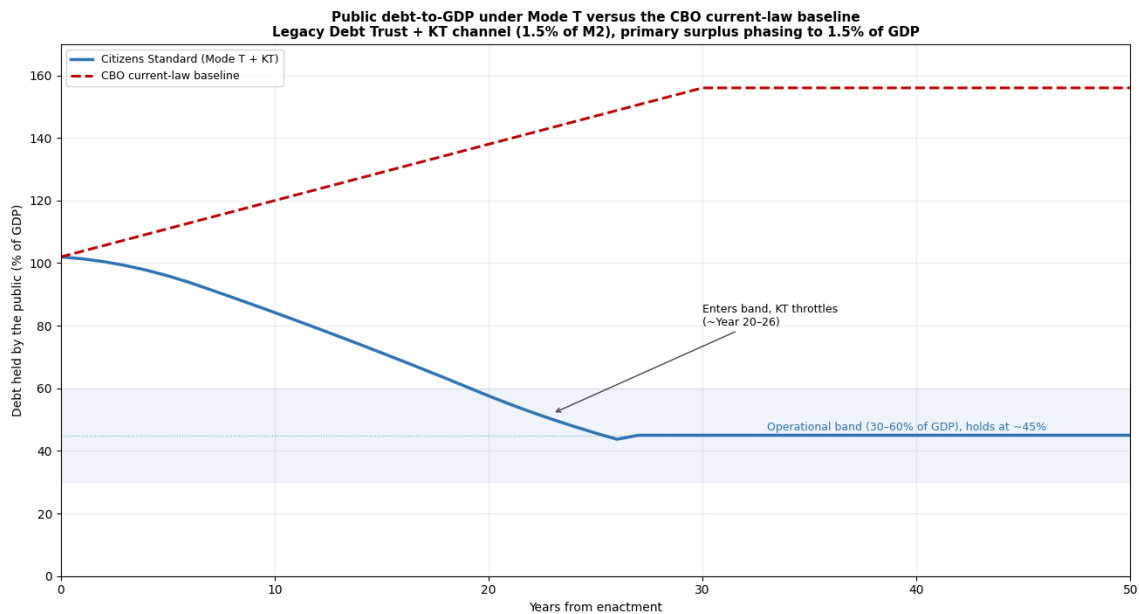


Figure 1. Public debt-to-GDP under Mode T versus the CBO current-law baseline. With the Legacy Debt Trust refinancing the \$31.4 trillion held by the public, the KT channel issuing ~1.5% of M2 to redemption, and a primary surplus phasing in to 1.5% of GDP, public debt falls from 102% at enactment into a moderate operational band of approximately 30 to 60 percent of GDP (central path ~45%, reached by Year 26), where it stabilizes rather than continuing toward zero — against the CBO’s 156% by 2055 under current law. KT self-throttles as the debt enters the band and goes dormant within it, leaving a standing stock that supplies the safe-asset benchmark while the freed seigniorage is routed to citizen Stable Floors.

Public debt-to-GDP declines from 102% at enactment to approximately 58% by Year 20 and stabilizes within a moderate operational band of approximately 30 to 60 percent of GDP (central path ~45%, reached by Year 26) — materially better than the CBO’s March 2025 projection of 156% by 2055 under current law. Thereafter the system lands automatically in Mode T-stable, a

permanent steady state that is Mode Ω , holding the price level stable for any transaction ratio (coinciding with the fixed Mode B split at the one-half ratio, and given by the Ω solve otherwise), in which KT remains available as a symmetric open-market instrument for two-sided monetary control.

It is worth locating the KT mechanism against the historical record of sovereign-debt reduction, because it does not belong to any of its standard categories. Large debt burdens have historically been reduced through five channels: real growth that outpaces the debt, sustained primary surpluses, outright inflation that erodes the real value of nominal claims, financial repression that holds nominal rates below inflation so that bondholders are taxed implicitly (the channel Reinhart and Sbrancia (2015) identify as the dominant force behind postwar advanced-economy deleveraging), and default or restructuring. KT is none of these. It retires the public debt through seigniorage-funded redemption — new money created under the same constitutional rule that funds citizen Stable Floors, used to redeem maturing bonds — while a price-level throttle automatically contracts issuance if consumer inflation rises, so the mechanism cannot become the inflationary debt erosion it superficially resembles. This distinguishes it equally from the overt monetary financing debated by Turner (2016), which lacks a binding non-inflation constraint, and from financial repression, which achieves reduction by transferring the cost to savers rather than by routing the seigniorage to citizens. The transition's contribution on this point is to show that debt retirement and citizen capital formation can be funded by the same rule without the consumer-price consequences that each of the five historical channels carries.

4.1.1 The Operational Band: A Debt-Sustainability Analysis of the Endpoint

The transition retires the public debt as a fiscal burden; it does not follow that the debt should be retired as close to zero as the mechanism allows. The endpoint is a design choice, and it is set by a welfare criterion rather than by the impulse to minimise the stock. We pose it as a stochastic debt-sustainability analysis, calibrated to the empirical literature on the interest-growth differential, and the answer is a moderate band of approximately 30 to 60 percent of GDP rather than a low floor.

The criterion balances two opposing considerations on a common welfare scale. On one side, standing sovereign debt is near self-financing whenever the safe interest rate sits below the growth rate, which the long-run record shows is the historical norm rather than the exception (Blanchard 2019; Mauro and Zhou 2020): the growth snowball lowers the debt ratio on its own, so seigniorage spent retiring the stock below the level required for sustainability is seigniorage not delivered to citizen Stable Floors. Aggressive retirement therefore carries a large and otherwise invisible opportunity cost, measured in the citizen floors forgone. On the other side, the probability and severity of an adverse reversal in the interest-growth differential rise with the debt level itself: Lian, Presbitero and Wiriadinata (2020) document that as debt rises from 40 to 120 percent of GDP the ninetieth percentile of the five-year-average differential rises from roughly zero to about two percentage points, and the likelihood of a reversal from a favourable to an unfavourable regime rises from about one quarter below median debt to more than three quarters in the top quartile. This crisis channel caps the band from above.

We represent the differential by a compound-jump process: a tightly negative baseline regime, punctuated by rare positive jumps whose magnitude generates the severity of adverse episodes and whose persistence generates sustained reversals, with both the

jump frequency and persistence rising in the debt ratio. A Gaussian regime-switching specification cannot reproduce the data because it ties the magnitude of the upper tail to the frequency of unfavourable years through a single variance; the jump specification separates the two, and reproduces the full set of severity quantiles and the reversal-frequency gradient of Lian et al., together with a held-out moment, the share of years with a negative differential, that it was not fit to. Maximising expected welfare, citizen seigniorage net of an expected-crisis term over the resulting debt-path distribution, yields an interior optimum: a moderate band rather than either corner. The optimum is a shallow maximum, so the defensible object is a range of roughly 30 to 60 percent of GDP rather than a single point.

The one parameter the band depends on materially is the issuer's exposure to debt-driven crisis risk. Here the Citizens Standard occupies a structurally favourable position. The crisis channel that the analysis prices is, in substantial part, the self-fulfilling liquidity dynamic that afflicts sovereigns issuing in a currency they do not control; it does not arise for a sovereign with its own currency and its own monetary authority, whose central bank can always provide the liquidity to meet a maturing obligation in nominal terms (De Grauwe 2011, on the contrast between standalone monetary sovereigns and members of a monetary union). The Citizens Standard issues in its own unit and operates its own issuance authority, so it has this fragility immunity by construction, which places it toward the damped end of the crisis-risk range and the upper end of the band. It does not, at inception, possess the separate property of an established international reserve currency, the active flight-to-quality in which global investors compress yields during stress; that is earned over decades and is not assumed here. Two honest qualifications follow. First, monetary sovereignty removes the illiquidity and self-fulfilling-panic channel, not the fundamental-solvency channel: a genuine over-issuance or inflation crisis remains possible, which is why the band has a ceiling at all and should not be set high. Second, the band is stated as a range precisely because its centre slides with this single assumption, toward the lower thirties if the issuer's safe-haven standing is treated conservatively and toward sixty if treated strongly; the moderate-band conclusion itself is invariant to the choice.

The operational implication for the KT channel is a change of target, not of mechanism. KT retires debt at full rate until the stock enters the band, then throttles to hold position, so that the growth-matched seigniorage no longer needed for retirement is routed to citizen Stable Floors. The form in which the standing stock is then held, non-interest-bearing by default, is the subject of the next section. The full analysis, calibration, verification against the published moments, and the welfare optimisation, is released as the `cs_debt_band` modules in the replication package.

4.1.2 The Form of the Standing Stock: Non-Interest-Bearing by Default

The band fixes how much sovereign liability to hold; a second choice fixes its form. The standing stock should be held mostly non-interest-bearing, and the reason is a clean separation of what a safe public liability provides from what a coupon costs. A safe public liability earns a convenience yield: holders accept a financial return below the otherwise-fair rate because they value its safety, liquidity, and collateral function (Krishnamurthy and Vissing-Jorgensen 2012 estimate this premium at roughly seventy basis points historically, compressed toward the mid-thirties at today's high supply; St. Louis Fed 2025). That convenience is a function of the stock, and the issuer earns it as cheaper-than-fair funding whether or not the instrument carries a coupon. The coupon itself buys nothing the convenience does not already supply: it is a transfer from the public to the holder.

The magnitudes make the point. On a stock at the band's central path, approximately 45 percent of GDP, a conventional coupon of three percent transfers on the order of 1.35 percent of GDP per year to holders, a class that skews wealthy, for no public benefit the convenience value does not already deliver. Holding the same stock at the non-interest-bearing limit removes that transfer and captures the convenience as public seigniorage instead, routed to the same pool that funds citizen Stable Floors. The instrument at that limit is economically money: a par-fixed, default-free public liability whose holders are compensated by its convenience rather than by interest. This is why the standing stock is not a fiscal burden in the form recommended here; it is a monetary instrument that pays the public rather than the bondholder.

The non-interest-bearing limit gives up one thing, and the paper states it plainly: a traded benchmark yield curve. The wider financial system prices corporate bonds, mortgages, and derivatives off a coupon sovereign curve, and a fully non-interest-bearing stock supplies safety and collateral but not that pricing benchmark, which requires a traded, conventionally-priced instrument across maturities. The resolution is a thin tranche: retain a small coupon stock sufficient to anchor the benchmark curve while holding the bulk of the standing stock non-interest-bearing. The transfer is then confined to the tranche actually needed for the benchmark function, and the rest of the convenience accrues to the public.

One conditional is worth naming because it sets the long-run direction. The convenience yield is not a constant; it shrinks as engineered-stable money absorbs safe-asset demand. The more completely the framework's own money becomes the economy's safe store of value, the less a separate safe bond is demanded, the smaller the convenience yield, and the further the standing stock can shrink toward money itself. In the limit the standing stock and the circulating money converge: a moderate stock of non-interest-bearing public liability is the same object whether it is reached by retiring debt to a zero-coupon band or by issuing debt-free public money directly. The form recommended here is therefore not only the welfare-maximising choice within the framework; it is the point at which this framework and a debt-free public-money system describe the same steady state.

4.2 Banking Separation and Credit Stability

The US banking system currently holds approximately \$18 trillion in deposits against roughly 10 percent in reserves — a gap of approximately \$16.2 trillion that cannot be closed without collapsing the credit system. The annual pace of reserve conversion over 20 years — approximately 4.5 percentage points per year — puts approximately \$810 billion in annual credit at risk of contraction.

The Transition Lending Facility offsets this by channeling a bounded share ($\gamma \leq 0.30$) of K2 seigniorage — the K2 line itself growing from approximately \$219 billion annually at Phase 3 entry to approximately \$438 billion at full Mode B — into replacement lending at zero spread for community banking, agricultural lending, and small business finance. TLF coverage of annual credit at risk ranges from approximately 12 to 38 percent at full-rate K2, with KT's sovereign-money injection (approximately 41 to 59 percent) supplying the larger complementary share. The TLF is a targeted replacement of the most vulnerable channels rather than a complete offset; the residual credit contraction beyond its coverage is addressed by a second, conditional instrument. Where the TLF replaces vulnerable lending on a scheduled basis, the conditional KI_T damper (macro model, Neo-Solon 2026e, Section 4.7) absorbs the residual contraction as a price-path feedback rule: it is a non-binding-by-default channel, in force throughout the transition but slack while the realized price level holds to target, becoming binding only when the residual contraction

pulls the path below target. KI_T issues additive, net-new spendable money equally per citizen — never a diversion from K1 or K2, so citizen Stable Floor flows remain untouched throughout — its intensity is a ratifiable parameter, and it is withdrawn from the configuration once banking separation completes. Its stability is inherited from the path-targeting convergence result (Neo-Solon 2026e, Proposition 3): with a closure gain no greater than the steady-state rule, it converges monotonically and cannot oscillate. Large corporate credit markets have alternative financing channels that do not depend on bank deposit funding; the TLF and the conditional KI_T damper together address the bank-dependent residual.

Three conversion speed scenarios: a fast 10-year conversion carries high recession risk; a 20-year base case carries moderate recession risk manageable through the emergency toolkit; a 30-year slow conversion carries low recession risk. The operative lever is conversion pace rather than facility size: annual credit-at-risk equals the fixed reserve gap divided by the conversion window, so a longer conversion lowers the per-year gap toward KT's roughly fixed injection capacity and raises KT's coverage fraction — closing the gap without enlarging the TLF, whose $\gamma \leq 0.30$ cap protects the Stable Floor flow. The 20-year base case is recommended. The economic literature on the Chicago Plan (Benes and Kumhof, 2012) suggests credit reduction need not be recessionary if offset by sovereign money creation, but the offset is partial, not complete. A society choosing this transition accepts tighter credit conditions in exchange for a payment system that cannot fail when the credit system does.

4.3 Equity Valuation Effects of Universal Ownership Flows

At full Mode B, K1 and the 60 percent of K2 that funds the locked floor direct approximately \$272 billion annually into total-market equity purchases — the remaining 40 percent of K2 is paid to citizens as the standing consumer dividend — approximately 0.39 percent of the US equity market's approximately \$69 trillion 2025 capitalization. For context, annual 401(k) and IRA contributions currently run approximately \$700 billion (1.01 percent of market cap); Federal Reserve QE peaked at approximately \$1.44 trillion annually in 2020–21 (2.09 percent); share buybacks in 2024 ran approximately \$900 billion (1.30 percent). The Citizens Standard's full-Mode-B flow is approximately 65 percent of current 401(k)/IRA flows and remains well below both QE and buyback magnitudes.

Using the price-impact literature's multiplier range of 1x to 3x for diversified systematic flows, the \$272 billion annual flow implies annual market capitalization increase of \$272 billion to \$816 billion — 0.39 to 1.18 percent of market cap per year. The transition-specific return compression — the portion that reverts once KT sunsets — is driven by the KT bond-rotation flow and is approximately 0.4 to 0.6 percentage points per year, within the bootstrap confidence intervals reported in the empirical paper (Neo-Solon, 2026b). The permanent Stable Floor flow's price impact is a steady-state feature, not a transition cost, and accrues substantially back to citizens who hold the appreciating broad-market equity. The equity valuation question qualifies the framework's projections modestly downward, in a direction already captured by the empirical paper's low-band realizable-return case (3.30 percent real).

The 401(k) expansion from 1980 to 2000 provides the most direct historical analogue: assets grew from approximately zero to \$3.5 trillion — an average inflow of approximately \$175 billion annually, comparable to full-Mode-B flows. Attribution of the CAPE expansion over this period to 401(k) flows specifically is partial at most, given concurrent falling rates, the technology productivity boom, and strong earnings growth.

4.4 Constitutional Lock Timing

Three conditions must be jointly satisfied for constitutional lock to be credible: sufficient electoral constituency with direct stake, account balances visible and politically salient, and survival of at least one major market cycle without emergency legislative intervention.

Electoral analysis shows that at Year 25, approximately 52 to 83 million citizens hold Stable Floor accounts — Phase 1 birth cohorts reaching voting age (approximately 10 million) plus Phase 2 voluntary opt-ins (42 to 74 million at 20 to 35 percent participation). At Year 35, Phase 1 cohorts carry observable balances of approximately \$55,000 to \$90,000 in 2025 real dollars.

The **optimal lock window is Year 35 to 45**, coinciding with the Phase 4 to Phase 5 transition: full Mode T is operational, the Federal Reserve conversion is underway, and the oldest Phase 1 cohorts are within 20 to 30 years of retirement with observable compounding trajectories. The hard outer boundary is Year 50 — constitutional protection must precede the first Phase 1 retirement claims.

Social Security's political durability emerged approximately 25 to 30 years after its 1935 enactment. Medicare achieved durability approximately 15 to 20 years after 1965. Both timelines are consistent with the Citizens Standard's Year 25 to 35 credibility window.

5. Failure Mode Analysis

Phase 1 — Political reversal. Constitutional lock on existing balances at inception converts reversal risk from account elimination to deposit cessation — existing compounding continues even if new deposits stop.

Phase 2 — Inadequate voluntary participation. Automatic enrollment with opt-out rather than voluntary opt-in produces participation rates of 85 to 90 percent versus approximately 49 percent under opt-in, per Madrian and Shea (2001). This converts participation risk from structural to residual.

Phase 3 — Credit contraction recession. The emergency toolkit provides approximately \$3.0 to \$3.5 trillion in combined first-year demand-collapse response capacity, raised to approximately \$3.7 to \$4.1 trillion with the addition of Tool 15 (M2 Contraction Floor, 3.0 percent of M2 annually), which targets the Fisher debt-deflation cascade specifically — triggered when M2 contracts more than 5 percent over any rolling 12-month window. The dynamic cascade model in the replication package indicates the full toolkit offsets roughly a third of a Depression-magnitude M2 contraction and over half of a 2008-magnitude one, with Tool 15 contributing approximately 3 to 4 percentage points in acute scenarios and its 18-month sunset bounding its contribution in prolonged contractions; the offset is partial, not complete, and the credit system still bears real stress while the payment system remains protected throughout. The Phase 3 tranche structure provides natural checkpoints — if credit contraction exceeds thresholds in Tranche 1, Tranche 2 can be delayed without abandoning the overall conversion.

Phase 4 — Federal Reserve institutional resistance. Functional transfer rather than institutional elimination — payment system functions transfer to the FDCA, open market operations cease, supervisory role is bounded — reduces resistance and preserves organizational knowledge.

Phase 5 — Supermajority threshold not reached. By Year 35, decades of compounding data, tens of millions of account holders with direct financial stakes, and observable balances make the supermajority case through demonstrated performance rather than projection.

6. Political Economy of the Migration Path

Phase 1 authorization requires only simple legislative majority. The framing of universal ownership — every American child receives an equity stake at birth — has no ideological owner. The Alaska Permanent Fund passed two-to-one because universal distribution creates broad-based support that cuts across political divisions. Phase 1's annual cost of approximately \$3 to \$23 billion is fiscally modest enough to survive budget negotiations.

The Ballot 2 coalition requires neutralizing three institutional actors. Commercial banks lose seigniorage revenue under the full architecture but Phase 1 and Phase 2 do not threaten their core lending functions. The Federal Reserve faces functional transformation but its staff and institutional knowledge transfer to the FDCA. The Treasury loses deficit monetization capacity, but under Mode T the KT channel services the legacy debt directly through bond redemption, so nominal revenue flows are maintained without Treasury deficit financing. These are sequencing decisions, not permanent concessions.

The generational transfer. The phase architecture's deepest political economy insight is that its primary beneficiaries — citizens born after Phase 1 launch — are not yet political actors when critical ratification decisions are made. The political coalition is built by the program's own operation: Phase 1 creates account holders with direct financial stakes, Phase 2 expands that constituency, and the compounding trajectory makes the program's value increasingly visible as the Ballot 2 window approaches.

7. International Transition Analogues

Chile: Price-indexed debt and the Unidad de Fomento (1967–present). Created by Decree 40 of the Chilean Ministry of Finance on January 20, 1967, the UF is a price-indexed unit of account adjusting daily with CPI, now used in virtually all Chilean mortgages, long-term contracts, and government bonds. Chile's experience demonstrates that price-indexed debt is operationally manageable at national scale, that the transition from nominal to indexed instruments can be accomplished gradually, and that price indexing reduces financial system fragility by eliminating inflation-driven redistribution between borrowers and lenders.

Norway: Sovereign wealth governance (1990–present). The Norwegian Government Pension Fund Global survived multiple changes of government, commodity price crashes, and two global financial crises without governance breakdown or politically motivated raiding. Three governance features apply directly: operational independence from the annual budget process; an ethical investment mandate enforced by an independent Council of Ethics; and full quarterly portfolio disclosure that builds public trust over time.

Estonia: Digital identity infrastructure (1991–present). Estonia verifies unique, living personhood for 1.3 million citizens using privacy-preserving cryptographic methods and a decentralized architecture. The architectural paper estimates a 3 to 5 year US

infrastructure development timeline. Estonia built its core system between 1996 and 2002 — approximately six years — and has operated it continuously since.

Singapore: Mandatory locked accumulation (1955–present). The CPF has operated across 70 years of dramatic change. Mandatory participation is sustainable when benefits are visible and individually owned; the architecture can expand its social function without losing structural integrity; and contribution rate adjustments are manageable when rules-based and transparent.

8. Discussion

What the transition analysis establishes. This paper has specified a credible migration path from the current monetary system to the Citizens Standard's constitutional architecture. The path is constructable from existing institutional materials. The debt transition problem, which initially appears to require an impossible fiscal adjustment, is resolved by separating the relevant debt — the \$31.4 trillion held by the public — from the intragovernmental balances that net out, transferring it to a Legacy Debt Trust that refinances but never expands the stock, and directing price-level-calibrated money creation through the transition-only KT channel to bond redemption. Because redemption is an asset swap absorbed by a reinvesting holder base, the mechanism retires debt while remaining consumer-price neutral and self-throttling on inflation. The recommended path produces public debt-to-GDP declining from 102 percent at enactment to approximately 58 percent by Year 20 and stabilizing within a moderate operational band of roughly 30 to 60 percent of GDP (central path approximately 45 percent, reached by Year 26) — retiring the debt as a fiscal burden while keeping a standing stock as the safe-asset benchmark and the base for the symmetric reverse-KT drain — materially better than the CBO's March 2025 projection of 156 percent by 2055 under current law, after which the system lands automatically in the permanent Mode T-stable steady state. The band rather than a low floor is the welfare-optimal endpoint: a stochastic debt-sustainability analysis calibrated to the interest-growth-differential literature shows that standing debt is near self-financing while the safe rate stays below the growth rate, so retiring below the band would forgo citizen seigniorage for no sustainability gain, while the band stays well clear of the debt levels at which crisis risk rises.

Supplementary replication. A debt-dynamics check on this retirement path, run on current fiscal data (debt held by the public anchored at ~100 percent of GDP), characterizes both the descent and the choice of endpoint. The descent into the band by approximately Year 20 is driven primarily by nominal GDP growth expanding the denominator, with the KT channel (on the order of 1.5 percent of GDP per year) accelerating the path at the margin and guaranteeing that no new borrowing is issued; this is how postwar economies reduced their debt ratios, by growing out of the debt under price stability. The endpoint — a moderate band of approximately 30 to 60 percent of GDP rather than a low floor — is established by a stochastic debt-sustainability analysis calibrated to the published interest-growth-differential literature (Blanchard 2019; Mauro and Zhou 2020; Lian, Presbitero and Wiriadinata 2020). That analysis prices two opposing considerations on a common welfare scale: because standing debt is near self-financing while the safe rate sits below the growth rate, seigniorage spent retiring below the band is seigniorage not delivered to citizens (the snowball lowers the ratio for free), while the probability and severity of an adverse interest-growth reversal rise with the debt level, capping the band from above. The welfare-optimal stock

balances these, and the moderate-band conclusion is robust across the plausible range of the one swing parameter (the issuer's sensitivity to debt-driven crisis risk, which a sovereign-currency issuer with its own monetary authority structurally dampens; De Grauwe 2011). Runnable as the `cs_debt_band` modules in the replication package.

What the transition analysis does not establish. This paper does not demonstrate that the transition will succeed — only that it is technically feasible and each failure mode has a containment mechanism. The debt model assumes markets price indexed sovereign instruments at near-TIPS real yields during the transition window; modeling transition risk premiums under simultaneous banking restructuring and constitutional change is the primary unresolved quantitative question this paper leaves open. The treatment of existing Social Security obligations during the transition and the international reserve currency dynamics of a dollar transitioning to rules-based issuance require their own analyses and are candidates for subsequent papers in this series.

The comparative disruption argument. The Citizens Standard's migration path may appear disruptive relative to the current system's continuity. The comparison is misleading. The current system has already produced, within the past two decades: a global financial crisis requiring \$2.9 trillion in liquidity support; a monetary expansion of 40 percent over two years producing 9.1 percent inflation; annual net interest payments of \$970 billion with no ceiling in sight; and a retirement system leaving the median American with approximately \$95,000 in savings. The Citizens Standard's transition costs are real. They are smaller, more predictable, and more controllable than the disruptions the current system has already demonstrated it produces.

8.1 Social Security and the Citizens Standard: Consolidation as a Phase 4–5 Question

The Citizens Standard is designed to supplement Social Security, not replace it. The architectural paper (Neo-Solon, 2026a) states this explicitly, and the empirical paper (Neo-Solon, 2026b) includes Social Security benefits in every retirement income comparison. The two systems can coexist indefinitely — Social Security as the legacy defined-benefit floor funded by payroll tax, the Stable Floor as the new equity-based floor funded by seigniorage. Nothing in the Citizens Standard requires Social Security to be wound down.

A thoughtful society operating a mature Citizens Standard would, however, face a legitimate question in Phase 4 or Phase 5: once the Stable Floor has compounded for 40 to 50 years and citizens at retirement are holding real balances of \$0.5 to \$0.8 million, what is the ongoing justification for a parallel payroll-tax-funded defined-benefit system? Social Security at that point would represent a meaningful fiscal drag — approximately \$1.3 trillion in annual outlays and a combined employer-employee payroll tax of 12.4 percent on all earned income — on a system that no longer structurally needs it. The political pressure to consolidate would be real, and ignoring this question in a transition paper would leave an obvious gap.

The honest position is this: consolidation is not architecturally required but is fiscally prudent at maturity, and the framework should address it as a Phase 4–5 option rather than either mandating it or pretending it will not arise.

Why consolidation makes structural sense at maturity

The payroll tax that funds Social Security is a labor tax. It raises the cost of employment and suppresses wages for every worker below the taxable maximum — currently \$184,500 — regardless of their Stable Floor balance. In a Phase 5 world where every citizen has a constitutionally guaranteed capital stake compounding from birth, the payroll tax is financing a benefit that the Stable Floor already provides more effectively. The net effect is a deadweight loss: workers bear a 12.4 percent tax on earned income to fund a retirement floor that their Stable Floor makes redundant for the median citizen and that already sits below the Stable Floor for the modeled cohorts on the realizable basis.

A second structural argument concerns the distributional properties of the two systems. Social Security replaces a fixed percentage of pre-retirement earnings, meaning higher earners receive larger nominal benefits. The Citizens Standard distributes equal per-citizen K1 and K2 deposits regardless of earnings history. As the Citizens Standard matures, the distributional logic of Social Security — calibrated to earnings rather than citizenship — becomes increasingly in tension with the Citizens Standard's constitutional commitment to equal per-citizen issuance.

How consolidation would work — a Phase 4–5 framework

Three constraints must govern any consolidation to protect current and near-term recipients. First, no reduction for citizens already receiving benefits or within 15 years of eligibility: these citizens' working lives predate the Citizens Standard and their Stable Floors reflect incomplete accumulation windows. Second, cohort-by-cohort phase-down for younger workers: citizens born after Phase 1 launch accumulate Stable Floor balances from birth; the payroll tax rate for these cohorts can be reduced gradually as their Stable Floor balances grow toward retirement, with full elimination only when Stable Floor projections reliably exceed Social Security replacement value. Third, a constitutional floor guarantee: any consolidation legislation must include a constitutional provision ensuring that Stable Floor withdrawals at age 65 cannot, in combination with any reduced Social Security benefit, produce total retirement income below the poverty line. This is the hard floor the consolidation must protect.

The payroll tax revenue freed by phase-down provides a material fiscal dividend. At 12.4 percent of approximately \$10 trillion in covered wages, full elimination represents approximately \$1.24 trillion in annual tax relief — effectively returning the cost of the old retirement system to workers and employers as wages and investment. This is not a fiscal windfall; it is the structural consequence of having funded the replacement system through seigniorage rather than taxation for four decades.

What this paper establishes and does not establish

This section does not recommend consolidation. It establishes that consolidation is a coherent Phase 4–5 option, that it is structurally motivated by the maturity of the Citizens Standard rather than by ideology, and that it can be executed safely with appropriate protections for current and near-term recipients. A society that reaches Phase 5 and decides to maintain both systems in parallel is making a legitimate constitutional choice. A society that decides to consolidate is also making a legitimate constitutional choice — one this paper argues deserves a framework rather than silence. The Citizens Standard's core commitment to citizen sovereignty over monetary and fiscal architecture means this question, like Mode selection, is properly the citizens' to answer through

supermajority process at the appropriate phase — not something a transition paper should resolve on their behalf.

9. Conclusion

The Citizens Standard can begin today. Not as a monetary revolution — as a sovereign wealth migration. Every institutional element required for Phase 1 exists: custodial infrastructure for equity accounts, total-market index funds with sub-0.05 percent fee structures, Social Security Administration identity verification for newborn citizens, and legislative precedent for universal distribution programs.

The migration from Phase 1 to full constitutional architecture is a 40 to 60 year process. It requires the Legacy Debt Trust to refinance the \$31.4 trillion of public debt without permitting new government borrowing, and the transition-only KT channel to direct price-level-calibrated money creation to bond redemption — retiring the public debt into a moderate operational band of approximately 30 to 60 percent of GDP while citizen Stable Floor accumulation continues uninterrupted, then landing automatically in the permanent Mode T-stable steady state, where KT persists as a symmetric open-market tool. It requires phased banking separation over 20 years with the Transition Lending Facility absorbing the most vulnerable credit channels. It requires a two-ballot ratification sequence that builds democratic legitimacy through demonstrated performance. And it requires constitutional lock at the optimal window of Year 35 to 45 — before the first Phase 1 cohort reaches retirement age, while the political coalition built by three decades of compounding is at its peak.

What this paper establishes is that none of these obstacles is insurmountable, and that the phase architecture addresses each with mechanisms drawn from existing institutional precedent. The Alaska Permanent Fund proved that universal distribution of seigniorage flows is politically durable. Norway proved that sovereign wealth governance at national scale is administratively manageable. Chile proved that price-indexed debt works. Estonia proved that privacy-preserving digital identity is buildable. Singapore proved that mandatory locked accumulation is sustainable across generations. The Citizens Standard's transition does not require inventing new institutions — it requires sequencing existing ones correctly.

The architectural paper (Neo-Solon, 2026a) specified the destination. The empirical paper (Neo-Solon, 2026b) demonstrated that the destination's retirement-security claims survive historical stress-testing. This paper has specified the path. The research program's next questions are the ones this paper defers: the interaction between the Citizens Standard transition and Social Security's long-run solvency, the international reserve currency dynamics of a rules-based dollar, and the broader political economy of constitutional monetary reform in comparative perspective.

References

- Benartzi, Shlomo, and Richard H. Thaler. "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving." *Journal of Political Economy* 112, no. S1 (2004): S164–S187.
- Benes, Jaromir, and Michael Kumhof. "The Chicago Plan Revisited." IMF Working Paper WP/12/202, 2012.
- Blanchard, Olivier. "Public Debt and Low Interest Rates." *American Economic Review* 109, no. 4 (2019): 1197–1229.
- De Grauwe, Paul. "The Governance of a Fragile Eurozone." CEPS Working Document No. 346, May 2011.
- Lian, Weicheng, Andrea F. Presbitero, and Ursula Wiriadinata. "Public Debt and $r - g$ at Risk." IMF Working Paper WP/20/137, 2020.
- Mauro, Paolo, and Jing Zhou. " r minus g negative: Can We Sleep More Soundly?" IMF Working Paper WP/20/52, 2020 (also IMF Economic Review, 2021).
- Bureau of Economic Analysis. National Income and Product Accounts Tables (GDPA series). US Department of Commerce, various years.
- Bureau of Labor Statistics. Consumer Price Index Historical Data (CPI-U, Dec-Dec). US Department of Labor, various years.
- Congressional Budget Office. The Long-Term Budget Outlook: 2025 to 2055. March 2025.
- Damodaran, Aswath. "Historical Returns on Stocks, Bonds and Bills: 1928–2025." NYU Stern, January 2026.
- Dimson, Elroy, Paul Marsh, and Mike Staunton. *Triumph of the Optimists: 101 Years of Global Investment Returns*. Princeton University Press, 2002.
- Federal Reserve Bank of St. Louis. M2 Money Supply (M2SL). FRED, 2025.
- Federal Reserve Board. Survey of Consumer Finances, 2022.
- Friedman, Milton. *A Program for Monetary Stability*. Fordham University Press, 1960.
- Madrian, Brigitte C., and Dennis F. Shea. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics* 116, no. 4 (2001): 1149–1187.
- Neo-Solon (2026a). *The Citizens Standard: One Model, Many Systems — A Constitutional Monetary Architecture*. SSRN Working Paper 6702518.
- Neo-Solon (2026b). *The Citizens Standard: A Historical Counterfactual — Empirical Analysis of an Alternative US Monetary Architecture, 1960–2055*. SSRN Working Paper 6735078.
- Neo-Solon (2026d). *The Citizens Standard: A Statutory Implementation Pathway*. SSRN Working Paper 6873798.
- Neo-Solon (2026e). *The Citizens Standard: A Macroeconomic Model of a Two-Circuit Monetary System*. SSRN Working Paper 6939418.

Reinhart, Carmen M., and M. Belen Sbrancia. "The Liquidation of Government Debt." *Economic Policy* 30, no. 82 (2015): 291–333.

Shiller, Robert J. *Irrational Exuberance*, 3rd edition. Princeton University Press, 2015.

Siegel, Jeremy J. *Stocks for the Long Run*, 5th edition. McGraw-Hill, 2014.

Social Security Administration. *The 2025 Annual Report of the Board of Trustees*. March 2025.

US Census Bureau. *National Population Estimates, Vintage 2025*. January 2026.

US Department of the Treasury. *Monthly Statement of the Public Debt*. March 2026.

Turner, Adair. *Between Debt and the Devil: Money, Credit, and Fixing Global Finance*. Princeton University Press, 2016.

Vanguard Group. *How America Saves 2025*. Vanguard Research, 2025.

Technical Appendix

The Citizens Standard: Transition Architecture and Migration Mechanics. Neo-Solon (2026) · Working Paper · Companion to SSRN 6702518, 6735078, 6810741, and the statutory implementation paper (2026d).

This appendix provides the mathematical specification of the four quantitative models in the paper: the sovereign debt transition model (Section 4.1), the banking separation and credit stability model (Section 4.2), the equity valuation flow model (Section 4.3), and the constitutional lock credibility model (Section 4.4). The debt model has been revised to reflect the public-debt-only framing, the Legacy Debt Trust, and the transition-only KT channel operating under Mode T. All parameter values are stated with sources. Python replication code is released as a supplementary file on SSRN.

A.1 Notation and Shared Variables

The following variables appear across two or more models. Model-specific variables are defined in their respective sections.

Symbol	Description	Value at launch	Source
t	Year index from launch (t=0 is launch year)	—	—
GDP(t)	Nominal US GDP in year t	\$30,762B	BEA GDPA Apr 2026
g_n(t)	Nominal GDP growth rate	1.8%/yr	Mode T price stability
g_r(t)	Real GDP growth rate	1.8%/yr	CBO LT Outlook 2025
pi(t)	CPI inflation rate	0.0%/yr	Mode T (price stable)
M2(t)	M2 money supply	\$22,366B	FRED M2SL
D_pub(t)	Federal debt held by the public	\$31,400B	US Treasury / CRFB 2026
D_intra	Intragovernmental debt (nets out)	\$7,600B	CRFB 2026
d(t)	Public-debt-to-GDP = $D_{pub}(t)/GDP(t)$	102%	Computed
r_legacy	Average legacy sovereign coupon	4.5%	US Treasury 2026
r_new	Coupon on refinanced instruments	~1.5%	Post-transition real yield
KT(t)	Transition debt-retirement issuance	~1.5% of M2	Mode T calibration
N(t)	US population	341.8M	Census Vintage 2025
MktCap(t)	Total US equity market cap	~\$69,000B	Wilshire 5000 full-cap, YE 2025 (\$69.1T)

A.2 Sovereign Debt Transition Model (Section 4.1)

The model rests on three corrections to a naive debt-stabilisation analysis: the relevant debt is the public debt, not the gross total; the Legacy Debt Trust resolves the rollover wall without permitting new government borrowing; and the transition-only KT channel retires debt through asset-swap money creation that is consumer-price neutral.

A.2.1 The Public-Debt Correction

Gross federal debt is approximately \$39.0 trillion, but \$7.6 trillion is intragovernmental — non-marketable special-issue securities held by the Social Security OASI fund, federal and military retirement funds, and Medicare Hospital Insurance. Per the CRFB, intragovernmental debt has no net effect on the government's overall finances; it does not trade and is redeemed internally. The Legacy Debt Trust therefore refinances only the public debt:

$$D_{\text{pub}}(0) = \$39.0\text{T} - \$7.6\text{T} = \$31.4\text{T} \quad (102\% \text{ of GDP})$$

The intragovernmental obligations are addressed by the Social Security consolidation question (Section 8.1), not by the Trust.

A.2.2 The Rollover Wall and the Legacy Debt Trust

With an average maturity of approximately six years, roughly one-third of the public debt — on the order of \$10–12 trillion — matures annually and must be refinanced. The Legacy Debt Trust is the sole entity permitted to refinance maturing pre-enactment securities; it may replace maturing instruments but never expand the stock. The government issues no new debt; the Trust winds the legacy stock down. The average coupon evolves as legacy instruments roll into refinanced ones:

$$\begin{aligned} r(t) &= w_L(t) \times r_{\text{legacy}} + [1 - w_L(t)] \times r_{\text{new}} \\ w_L(t) &= \max(0, 1 - t / T_{\text{roll}}), \quad T_{\text{roll}} = 6 \text{ years} \\ r(6) &= 3.0\% \text{ nominal } (\sim 1.5\% \text{ real}) \quad (\text{fully refinanced at the post-} \\ &\quad \text{transition yield}) \end{aligned}$$

A.2.3 The KT Channel

Under Mode T, the transition-only KT channel issues money calibrated to a price-level path, directed to the Trust for bond redemption rather than to citizens. Calibration is approximately 1.5 percent of M2 per year:

$$KT(t) = \text{kappa} \times M2(t), \quad \text{kappa} \approx 0.015$$

Because redemption is an asset swap absorbed by a holder base dominated by foreign central banks, pension funds, and the central bank (blended marginal propensity to consume ~2.5%), the consumer-price impact is negligible (+0.04pp expected, +0.16pp under a pessimistic 15% MPC). The channel is self-throttling: calibration to a price-level path reduces issuance automatically if consumer inflation rises. KT retires the public debt into a moderate operational band of approximately 30 to 60 percent of GDP (central path ~45 percent, reached by ~Year 26), eliminating the debt as a fiscal burden while retaining a standing stock that supplies the safe-asset benchmark and the base for symmetric reverse-KT operations. The band is the welfare-optimal endpoint of the debt-sustainability analysis (Section 4.1.1): standing debt is near self-financing at $r < g$, so retiring below the band forgoes citizen seigniorage for no sustainability gain, while crisis risk rises only at debt levels well above the band. What remains on the gross books is the ~\$7.6T

intragovernmental debt (~8% of grown GDP), handled by SS consolidation (Section 8.1). In the steady state KT is dormant but reverses — issuing bills above the band to drain transactional money — when consumer inflation runs above target, giving Mode T-stable two-sided monetary control.

A.2.4 Debt Dynamics and Trajectory

The debt stock evolves as maturing principal is retired by the primary surplus plus KT redemption:

$$D(t+1) = D(t) + r(t) \times D(t) - s(t) \times \text{GDP}(t) - \text{KT}(t)$$

with the primary surplus $s(t)$ phasing in from 0 to 1.5 percent of GDP over twenty-five years. The resulting public-debt-to-GDP trajectory:

Year	r(t)	KT (\$B)	Surplus (\$B)	D _{pub} (\$T)	D/GDP
0	4.50%	336	0	31.4	102%
10	3.00%	499	274	38.4	84%
20	3.00%	620	810	38.9	58%
30	3.00%	0	1498	45.0	45%
40	3.00%	0	2218	66.5	45%
45	3.00%	0	2,699	81.0	45%
CBO 2055	—	—	—	—	156%

Table A.2. Public-debt-to-GDP trajectory under Mode T. Legacy Debt Trust refinances \$31.4T public debt; KT directs ~1.5% of M2 to redemption; primary surplus phases to 1.5% of GDP over 25 years. Public debt retired into a moderate operational band of approximately 30 to 60 percent of GDP (central path ~45 percent, reached by Year 26), where the ratio holds while the nominal stock grows with GDP. Cumulative KT issuance ~\$12T; remainder of the descent delivered by surplus and nominal growth. The directional finding (band reached and held vs CBO 156% by 2055) is robust to reasonable variation in growth and surplus assumptions.

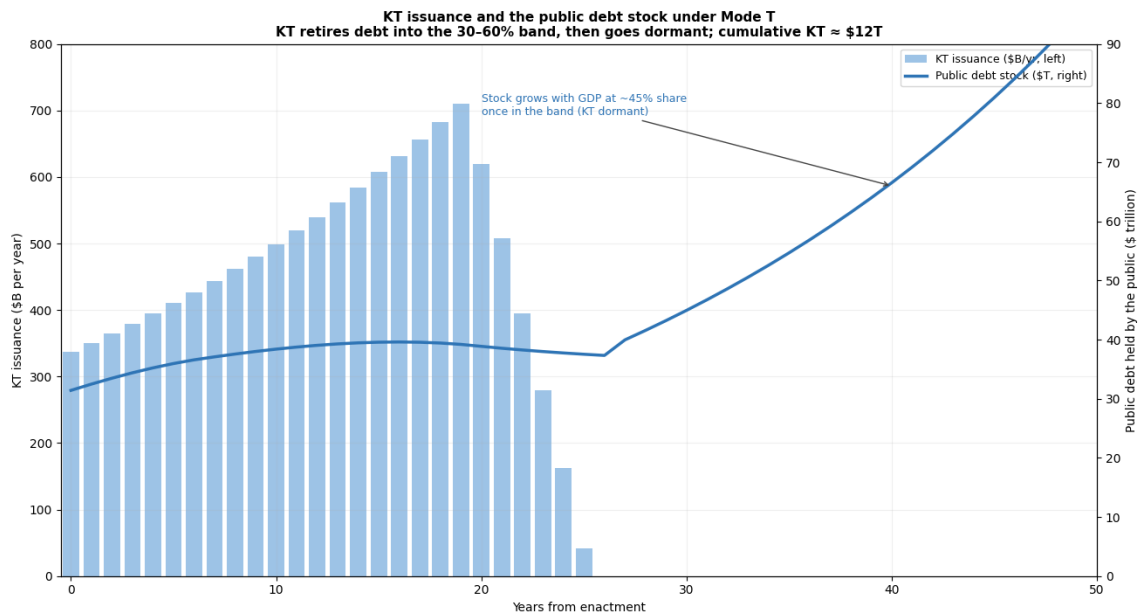


Figure A.1. KT issuance and the declining public debt stock. The KT channel issues ~1.5% of M2 per year to bond redemption (bars), retiring the public debt (line) from \$31.4 trillion into the 30-to-60-percent-of-GDP band. As the debt enters the band, KT self-throttles and goes dormant; the retained stock thereafter grows with GDP at a roughly constant ~45% share, held as the safe-asset

benchmark rather than as a fiscal burden. Cumulative KT issuance over the transition is ~\$12 trillion.

A.3 Banking Separation and Credit Stability Model (Section 4.2)

The banking separation model quantifies credit-contraction risk as fractional-reserve banks migrate to 100% reserve backing, and the offsetting capacity of the Transition Lending Facility (TLF). A synergy with the KT channel is noted at the end of this section.

A.3.1 The Reserve Gap

$$\text{Reserve gap} = D_{\text{dep}} \times (\text{RR}_{\text{target}} - \text{RR}_{\text{current}})$$

$$= \$18,000\text{B} \times (1.00 - 0.10) = \$16,200\text{B}$$

$$\text{Annual_CAR}(T) = \$16,200\text{B} / T$$

Scenario	Window	Ann. CAR	Recession risk	Rec.
Fast	10 yrs	\$1,620B/yr	High — exceeds TLF capacity	No
Base	20 yrs	\$810B/yr	Moderate — manageable	Yes
Slow	30 yrs	\$540B/yr	Low	Acceptable

A.3.2 The Transition Lending Facility

$$\text{TLF}(t) = \text{gamma}(t) \times K2_{\text{agg}}(t), \quad \text{gamma} \leq 0.30$$

At full Mode B, $K2_{\text{agg}} \approx \438B (the residual K1-funded line), so the $\text{gamma} \leq 0.30$ cap routes at most ~30 percent of K2 seigniorage to replacement lending and preserves the remainder for citizen Stable Floors. At Phase 3 entry (~Year 20), $\text{gamma} \approx 0.26$ directs ~\$57B to the TLF; the remainder flows to citizen Stable Floors. The TLF approaches its ceiling of ~\$131B/yr ($0.30 \times \438B) at full Mode B.

A.3.3 TLF Coverage and KT Synergy

Year	TLF capacity	Ann. CAR (20yr)	Coverage
20	~\$57B	\$810B	~7%
40	~\$131B	\$810B	~16%
35	~\$310B	\$810B	~38%

KT synergy. The full-reserve transition removes bank (inside) money-creation capacity, which is potentially deflationary. The KT channel injects sovereign (outside) money during the same window. The domestic portion of KT issuance offsets an estimated 41–59% of the annual credit-at-risk during the Years 20–40 overlap — a meaningful buffer on top of the TLF’s 12–38% coverage. The two mechanisms are complementary: KT provides outside-money creation precisely as full-reserve conversion removes inside-money creation.

Cascade dynamics. The static reserve-gap arithmetic above does not capture the Fisher debt-deflation feedback in which credit losses force deleveraging, deleveraging contracts term deposits and M2, and the resulting asset-price deflation triggers further losses. A dynamic cascade model (released in the replication package) traces this loop under loss shocks from moderate (2008-equivalent, 3% per year) to Depression-magnitude (6% per year). Its central findings: because the transaction pool is full-reserve and constitutionally protected, the maximum possible M2 contraction is bounded by the term-deposit share (60 percent of launch M2), unlike 1930–1933 when all deposits were simultaneously at risk; the full emergency toolkit, including Tool 15, offsets roughly a third of a Depression-magnitude contraction and over half of a 2008-magnitude one; and the failure mode is credit-system stress borne by explicit risk-bearers rather than payment-system collapse. The model informed the addition of Tool 15 to the emergency toolkit. The offset is partial,

not complete — the framework’s claim is a better failure mode and materially reduced magnitude, not the elimination of monetary contraction under severe sustained stress.

A.4 Equity Valuation Flow Model (Section 4.3)

A.4.1 Annual Flow Magnitude

$$\text{Flow}(0) = K1_agg + K2_agg = \$9B + 0.6 \times \$438B \approx \$272B/yr$$

As a share of ~\$69T market cap: ~0.39% annually from Stable Floor purchases alone.

A.4.2 Bond-Rotation Addition under KT

KT bond redemption returns cash to holders who reinvest. With sovereign instruments scarcer, an estimated 15–35% of proceeds (central estimate ~17%, since foreign central banks and the central bank — ~44% of holders combined — rotate almost nothing to equity) flows to equities. Combined Stable Floor plus bond-rotation demand peaks at ~0.75% of market cap annually, declining thereafter.

Programme	Flow/yr	% of cap	Notes
CS Stable Floor (K1+K2)	~\$272B	0.39%	Passive total-market
+ KT bond rotation (~17%)	~\$60B	0.09%	Asset swap reinvestment
401k + IRA contributions	~\$700B	1.01%	For comparison
Fed QE peak (2020–21)	~\$1,440B	2.09%	For comparison

A.4.3 Return Compression

$$\Delta_{PE} = \epsilon \times (\text{Flow} / \text{MktCap}), \quad \epsilon \approx 3$$

Combined flows imply a total forward return compression of approximately 0.4–0.6 percentage points during the transition, driven by the KT bond-rotation flow (which reverts once KT sunsets); the permanent Stable Floor flow is a steady-state feature already reflected in the empirical paper’s return assumptions, not a transition cost. This is bracketed by the low end of the empirical paper’s realizable-return band (3.30 percent real); the locked floor’s position relative to median actual wealth — roughly 0.81× to 0.94× on the realizable basis — is reduced modestly during the transition window but the structural advantage is preserved. The sensitivity of these results to the rotation fraction, and the historical evidence bounding that fraction, are examined in A.4.4.

A.4.4 Rotation Sensitivity and Historical Grounding

The rotation fraction is an input to the equity-valuation effect, not to the solvency arithmetic. A formal sensitivity sweep across the full 15–35% range confirms this separation. The debt-retirement trajectory — 102% of GDP at enactment into a moderate operational band of approximately 30 to 60 percent of GDP (central path ~45%) — is invariant to the rotation fraction, because KT redemption is an asset swap whose solvency mathematics do not depend on where the freed capital is subsequently allocated. KT’s consumer-price neutrality is likewise unaffected, since the marginal-propensity-to-consume analysis governs the consumption leakage, not the equity-versus-bond split of reinvested proceeds. What the rotation fraction does move is the magnitude of the transient equity effect: combined Stable Floor plus bond-rotation demand ranges from approximately 0.47% to 0.57% of market capitalization across the 15–35% band — a spread of about 0.10 percentage points — and forward return compression ranges from roughly +0.27 to +0.64 percentage points. Both are transient and revert once KT sunsets, and both remain below the scale of familiar systematic flows (retirement-account

contributions run near 1.1% of market capitalization annually; the peak of quantitative easing reached roughly 2.1%). Even stressing the fraction to 50% — far above any defensible holder-mandate decomposition — lifts combined demand only to approximately 0.90% of market capitalization. The result is robust by construction rather than by assumption.

Historical precedent bounds where the realistic fraction lies. Two large-scale sovereign-debt reductions are instructive. The United Kingdom reduced public debt from roughly 250% of GDP in 1946 to approximately 50% by the mid-1970s, and Canada reduced federal net debt from roughly 68% of GDP at its 1995 peak to approximately 50% by the end of the decade. In neither episode did the capital released by retiring the debt concentrate into equities and produce a destabilising asset-price event; it was reabsorbed across the full asset spectrum — other fixed income, savings instruments, and the real economy. In both cases the reduction was accomplished primarily by a negative growth-corrected interest rate — nominal output growth exceeding the average coupon — supported by primary surpluses, which is precisely the lever KT relies upon as legacy coupons roll to the lower post-transition real yield. The historical record therefore places the realistic rotation fraction at the low end of the assumed range, which makes the central ~17% estimate conservative rather than optimistic. The bottom-up holder-mandate decomposition and this historical bound are mutually reinforcing: both point to a modest, well-absorbed rotation rather than a concentrated equity surge.

A.5 Constitutional Lock Credibility Model (Section 4.4)

The lock becomes durable when three conditions hold simultaneously.

A.5.1 Three Conditions

C1 — Electoral majority of account holders ($\geq 130\text{M}$ of 260M voting-age adults). C2 — Survival of at least one complete major market cycle (~ 7 years; met by Year 10–15). C3 — Visible balances within 20 years of retirement ($\$25,000+$ real, politically legible).

Year	Phase 1 voters	Opt-ins (30%)	Total AH	vs 130M
25	9.6M	63M	72.6M	56%
35	23.3M	63M	86.3M	66%
40	30.1M	80M	110.1M	85%
45	37.0M	95M	132.0M	102%

A.5.2 Credibility Window

All three conditions converge at Year 38–45 (optimal lock window), with Year 50 the hard outer bound — constitutional lock must precede the first Phase 1 retirement claims at Year 65 by a sufficient margin. Social Security achieved durability ~ 25 – 30 years post-enactment; Medicare ~ 15 – 20 years; both consistent with the Year 35–45 window.

A.6 Full Parameter Table

Parameter	Value	Source
Nominal GDP at launch	\$30,762B	BEA GDPA Apr 2026
Public debt at launch	\$31,400B (102%)	US Treasury / CRFB 2026
Intragovernmental debt (nets out)	\$7,600B	CRFB 2026
Gross federal debt	\$39,000B	US Treasury 2026
CBO baseline D/GDP by 2055	156%	CBO LT Outlook Mar 2025
Mode T nominal GDP growth	1.8%/yr	Real growth; price stable
Legacy average coupon	4.5%	US Treasury 2026
Refinanced coupon (r_new)	~1.5%	Post-transition real yield
Treasury average maturity	~6 years	US Treasury 2026
KT calibration (kappa)	~1.5% of M2	Mode T design
Blended MPC (bond holders)	~2.5%	Holder-base weighted
KT operational band	~30–60% of GDP	Mode T design
Primary surplus (steady)	1.5% of GDP	Phases in over 25 yr
Total US bank deposits	\$18,000B	FDIC 2025 Q4
Reserve gap	\$16,200B	\$18T × 0.90
Base conversion window	20 years	Paper recommendation
Ann. credit at risk (base)	\$810B/yr	\$16.2T / 20
Equity rotation fraction	15–35% (~17%)	Holder-mandate decomposition
US equity market cap	~\$69,000B	Wilshire 5000 full-cap, YE 2025 (\$69.1T)
Return compression (transition)	0.4–0.6 pp	KT-rotation flow; epsilon≈3