THE CITIZENS STANDARD: A HISTORICAL COUNTERFACTUAL

Empirical Analysis of an Alternative US Monetary Architecture, 1960–2055

Neo-Solon

Neo-Solon@hotmail.com

WORKING PAPER — JUNE 2026

COMPANION PAPERS

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

The Citizens Standard: 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 Issuance Engine (Neo-Solon, 2026i) · SSRN 6973261

The Citizens Standard: Crisis Behaviour and Failure Modes (Neo-Solon, 2026l) · SSRN 6973358

The Citizens Standard: Distribution and Wealth Inequality (Neo-Solon, 2026n) · SSRN 6973342

Abstract

This paper provides an empirical reconstruction of the Citizens Standard framework against US historical data and a forward projection of its transition period, serving as the empirical companion to the architectural paper (Neo-Solon, 2026a), the transition paper (Neo-Solon, 2026c), and the macroeconomic model (Neo-Solon, 2026e). The architectural paper proposes a constitutional monetary framework whose Mode B configuration routes new money to citizens through a locked, individually owned equity stake — the Stable Floor — alongside a liquid growth dividend. Part I asks: applied to actual US economic data from 1960 to 2025, what would the framework have produced for representative citizens, and how does that compare to what citizens actually accumulated under the discretionary system?

The central methodological commitment of this version is that the floor is evaluated on a *realizable* basis. A universal program that deposits a share of national saving into citizen-owned equity cannot earn the historically realized price-taker return on equities indefinitely: at universal scale those deposits deepen the aggregate capital stock, and the return the floor can actually earn is the general-equilibrium marginal product of that deepened stock, not the marginal product an atomistic saver took as given. The macroeconomic model (Neo-Solon, 2026e, Section 6.7) derives this attenuated return; we adopt its Mode B central value of 4.26 percent real (band 3.30 to 5.03 percent) and compound every cohort's deposits at it. We also model Mode B at its 60/40 split: sixty percent of the growth-matched issuance builds the locked floor and forty percent is paid as a liquid citizen dividend (K3).

On this realizable basis the headline finding is not a multiple of the median but the *universalization* of it. Across four cohorts born 1960–1990, the locked Stable Floor reaches roughly $210,000 to $245,000 in 2025 real dollars — approximately the median American household's actual retirement wealth (0.81× for the earliest cohort rising to 0.94× for the latest), and only about 0.31× to 0.37× the *mean*. That gap between median-parity and a fraction of the mean is the result, stated precisely: the framework delivers the median saver's outcome to every citizen without reproducing the concentration embedded in the average. Counting the liquid K3 dividend the citizen also receives (about $27,000 to $32,000 over a working life, or more if voluntarily reinvested) raises the total value captured to 0.91×–1.07× the median. The comparison that matters is not floor-versus-median but a universal floor at roughly the median against a distribution in which the bottom half of households reaches retirement with almost nothing.

The result is deliberately independent of Social Security. We lead with a wealth-to-wealth comparison because Social Security, on the 2026 Trustees' projections, faces depletion of its retirement trust fund around 2032–2034 and an automatic benefit reduction of roughly 22 percent absent congressional action. The Stable Floor carries none of that fragility: it is funded by seigniorage rather than payroll taxes, it is a funded asset rather than a pay-as-you-go transfer, and it has no worker-to-beneficiary ratio to collapse. The floor is therefore presented as a hedge against the Social Security shortfall, not as a complement that depends on it.

A decomposition shows that, on the realizable basis, approximately 81 percent of the floor derives from compounding and 19 percent from the deposited principal — a more principal-weighted split than the price-taker reconstruction implied, because the attenuated return does less of the work. Stress tests substituting Depression-era and stagflation-era sequences during peak accumulation reduce the floor below the median for most cohorts; a Monte Carlo that recenters the historical return distribution on the realizable mean (preserving volatility and sequence structure) finds the floor below the median actual benchmark in roughly 29 to 54 percent of paths, with a typical (P50) outcome of 0.86× to 1.48× the median across cohorts. Counting the dividend, voluntarily reinvested at an after-tax rate, lifts every cohort above the median and reduces the below-median share to 6–35 percent. Part II projects the

transition cohorts and finds the debt-retirement transition near-neutral to citizen wealth, costing a temporary 2 to 11 percent return compression that declines across successive cohorts.

The paper concludes that, measured through its age-65 balances, the Citizens Standard functions as a structural floor on lifetime capital ownership rather than a retirement program or a means of beating the market. Its contribution is distributional, not a multiple: it makes the median household's capital position the universal minimum, self-funded and demographically immune. Modest in multiples, radical in distribution.

JEL classification: D31, E42, E51, E52, E58, N12

Keywords: Citizens Standard; capital ownership; wealth distribution; general-equilibrium returns; cohort analysis; Monte Carlo simulation; monetary history

Contents

1. Introduction

The Citizens Standard (Neo-Solon, 2026a) — hereafter the architectural paper — proposes a constitutional monetary framework built on a dual-circuit architecture. Its Mode B configuration issues new money through two ownership channels: K1 at citizenship events, calibrated at 2.5 percent of GDP per capita; and K2 annually, drawn from the real-growth-matched issuance line (real GDP growth times prior-year M2). Under Mode B that growth-matched line is split: sixty percent routes into locked, individually owned, total-market equity index accounts — the Stable Floor — that citizens cannot access until age 65, and forty percent is paid as a liquid citizen dividend (K3). The architectural paper makes empirical claims about Mode B that the present paper evaluates directly, on the realizable-return basis described below.

This version differs from a naïve reconstruction in one decisive respect. The framework is universal, so its returns cannot be those of a price-taking individual investor. An ordinary saver treats the equity return as exogenous — the marginal product of a capital stock whose size their own saving does not move. A program that channels a share of national saving into citizen-owned equity for everyone, permanently, is not a price-taker: at universal scale those deposits deepen the aggregate capital stock, and a deeper stock earns a lower marginal product. Compounding universal deposits at the historically realized equity return would therefore commit a fallacy of composition — it would assume every citizen can simultaneously earn the return that was available precisely because most citizens were not earning it. The companion macroeconomic model (Neo-Solon, 2026e, Section 6.7) solves for the return the floor can actually earn in general equilibrium: the marginal product of the deepened stock, which falls as the citizen capital share rises. For Mode B at its 60/40 split, that realizable return is approximately 4.26 percent real, with a band of 3.30 to 5.03 percent reflecting plausible capital-share and depreciation parameters. Every cohort balance in this paper is compounded at the realizable return, not the realized historical series.

The consequence is that the framework's empirical case is distributional rather than a matter of outsized returns. On the realizable basis the locked floor lands at approximately the median American's actual retirement wealth — not several times it — and well below the mean. That is the honest and, we argue, the more important result: a universal floor at the median, self-funded and constitutionally locked, in a system where the median is a level half the population never reaches and the mean is inflated by concentration at the top. The method is otherwise straightforward: take the architectural paper's K1 and K2 formulas, apply them to actual US economic data from 1960 forward, compound the resulting deposits at the realizable return to 2025 purchasing power, and compare to what the Federal Reserve's Survey of Consumer Finances shows Americans actually accumulated. A Monte Carlo extension (Section 6) recenters the historical joint distribution of inflation and growth, and the historical return distribution, on the realizable mean to characterize the full distribution of outcomes.

1.1 Positioning Relative to the Architectural and Transition Papers

This paper is a companion to, not a replacement of, the architectural paper. It takes the framework's architecture as given and evaluates its capital-ownership claims empirically. It does not re-derive the framework's monetary theory — the argument for citizen seigniorage, the K-formula design rationale, the constitutional governance architecture, or the general-equilibrium return result it imports from the macroeconomic model. Those claims are established in their own papers. The two strands are separable: the architectural and monetary claims stand or fall on their theoretical grounds; the capital-ownership claims stand or fall on what the data shows once the realizable return is imposed. See the architectural paper (Neo-Solon, 2026a, SSRN 6702518). The transition paper (Neo-Solon, 2026c, SSRN 6810741) specifies how the system migrates from the current arrangement, including the debt-

retirement mechanism whose effect on citizen outcomes is projected in Part II. A statutory implementation pathway is developed in the fourth paper (Neo-Solon, 2026d); the two-circuit macroeconomic model, the price-stability and bounded-compression propositions, and — critically for this paper — the general-equilibrium realizable-return result are developed in the fifth (Neo-Solon, 2026e). This paper thus serves as the empirical companion to both the architectural paper (Part I, the steady-state counterfactual on the realizable basis) and the transition paper (Part II, the forward transition projection).

2. Methodology

2.1 Scope and Design

The analysis covers Mode B exclusively, at its 60/40 split (sixty percent of the growth-matched line to the locked floor, forty percent to the liquid K3 dividend). Four cohorts are tracked: citizens born in 1960, 1970, 1980, and 1990. The 1960 cohort is fully retrospective (all 65 working-life years are observed in historical data). The 1990 cohort is substantially projected (30 of 65 years fall after 2025). Because returns are imposed on the realizable basis rather than read from the realized series, the distinction between observed and projected years matters for the deposit path (which depends on realized growth, inflation, and money supply) but not for the return applied to compounding, which is the realizable return throughout.

2.2 Data Sources

Six variables per year drive the counterfactual. All figures are deflated to December 2025 purchasing power using the BLS CPI-U Dec-Dec chain before compounding at the realizable real return throughout.

Variable	Source	Notes
M2 money supply	FRED M2SL	End-of-period (December) values, seasonally adjusted
Nominal GDP	BEA NIPA / FRED GDPA	Annual; 2025: $30,762.1B
Real GDP growth	BEA chain-weighted (A191RL1A225NBEA)	Year-over-year; drives the K2 issuance line
US population	Census Vintage 2025	Mid-year resident estimates
CPI-U Dec-Dec	BLS all items	2024: 2.9%, 2025: 2.7%; used for deflation only
S&P 500 nominal total return	Damodaran (NYU Stern), Jan 2026	Used to characterize volatility/sequence structure, not the return level
Realizable real return	Neo-Solon, 2026e, §6.7	Mode B GE return 4.26% (band 3.30–5.03%); applied to compounding

Table 1. Data inputs. Note the realized S&P 500 series enters only through the Monte Carlo, to characterize volatility and sequence structure; the level of the compounding return is the general-equilibrium realizable return, not the realized series.

2.3 K1, K2, and the 60/40 Split

Mode B K1 is calibrated at 2.5 percent of GDP per capita at each new-citizen event. The annual growth-matched line is the full positive real GDP growth rate times prior-year M2, with K1 funded first from that line and the remainder — the residual — split 60/40. Sixty percent of the residual builds the locked floor (K2); forty percent is the liquid citizen dividend (K3), paid as spendable cash and not

deposited into the locked account. Deposits are deflated to 2025 purchasing power using the CPI ratio before compounding at the realizable real return, ensuring all reported Stable Floor balances are directly comparable to 2025 retirement-account balances. **Calibration note.** The growth-matched line is the same one the system settles into after the debt-retirement transition (Mode T-stable); Mode B and Mode T-stable share the same steady-state issuance, and the 60/40 split is the Mode B allocation of it (architectural paper, Neo-Solon, 2026a). All deterministic and Monte Carlo results in this paper are generated from the authoritative replication engine at this calibration, with the realizable return imposed on compounding.

2.4 The Realizable Return and Its Band

For the level of the compounding return we do not use the realized S&P 500 series. The realized 1960–2025 geometric real return of approximately 6.5 percent is a price-taker outcome: it is the return atomistic investors earned on a capital stock whose size their saving did not materially move. Imposed universally, the framework deepens that stock, and the marginal product falls. The macroeconomic model (Neo-Solon, 2026e, Section 6.7) solves the resulting general-equilibrium fixed point and reports, for Mode B at its 60/40 capture, a realizable return of 4.26 percent real centrally, with a band of 3.30 percent (lower capital share, higher depreciation) to 5.03 percent (higher capital share, lower depreciation). We treat 4.26 percent as the central case and report the band as the primary sensitivity, replacing the price-taker “pessimistic / central / historical” equity scenarios of earlier reconstructions. The realized series is retained only in the Monte Carlo, where its volatility and serial-correlation structure are preserved but its mean is shifted to the realizable return.

2.5 Pre-Committed Reporting Standards

This paper pre-commits to reporting all findings including unfavorable ones. Any cohort or scenario where the floor falls below the median actual outcome is highlighted, not buried — and on the realizable basis several do. The decomposition of the floor into deposited principal and compounding is reported explicitly, with a frank account of what it does and does not establish (Sections 4.5 and 9.2). The liquid K3 dividend is reported as cash first, with voluntary reinvestment shown as an option at after-tax rates rather than folded into the headline. Non-survivor comparisons are reported in Section 5. Monte Carlo findings in Section 6 are reported in full, including the substantial probabilities that the floor falls below median actual outcomes under resampling.

3. Historical Reconstruction

3.1 Dataset Construction

We constructed an annual dataset spanning 1960 through 2025 (extended to 1929 for the Monte Carlo in Section 6). All reported Stable Floor balances are in December 2025 purchasing-power terms throughout. The deposit path is built from realized M2, nominal GDP, population, real GDP growth, and CPI; the compounding return is the realizable return of Section 2.4.

3.2 K1 and K2 Calibration in Historical Context

Mode B’s calibration produces dollar values that vary substantially across the historical period. Three illustrative years (K2 figures shown are the 60 percent floor share of the residual, per citizen):

Year	M2 ($B)	GDP per capita	K1 deposit (nom.)	K2 floor per cit. (nom.)	CPI ratio to 2025
1960	$312	$3,002	$75.04	$17.06	10.87×

Year	M2 ($B)	GDP per capita	K1 deposit (nom.)	K2 floor per cit. (nom.)	CPI ratio to 2025
1990	$3,285	$23,891	$597.27	$138.43	2.46×
2025	$22,366	$90,000	$2,250.01	$814.16	1.00×

Table 2. Illustrative Mode B calibration across the historical period. K2 floor per citizen is the sixty-percent floor share of the post-K1 residual. The order-of-magnitude difference between 1960 and 2025 deposits is the central methodological consequence of running the framework against historical data: the architectural paper's published figures assume launch at 2025 parameters, while this reconstruction inverts that timing.

3.3 The Return Series Used in Compounding

Compounding uses the realizable real return of 4.26 percent (central) throughout, not the realized S&P 500 series. This is the single most consequential difference between this reconstruction and a naïve one. A naïve reconstruction that compounded universal deposits at the realized ~6.5 percent geometric real return would overstate the floor roughly threefold for the earliest cohort, because it would assume every citizen simultaneously earns a return that existed only because most citizens did not participate. The realized series is highly variable — real losses exceeding 30 percent in 1974 and 2008, real gains exceeding 25 percent in several years — and that variability is carried into the Monte Carlo of Section 6 as volatility and sequence structure around the realizable mean. But the level at which deposits compound, in every deterministic result, is the general-equilibrium realizable return.

4. Cohort Outcomes

4.1 Headline Results

For each of four cohorts we computed the Mode B locked Stable Floor at age 65 in 2025 real dollars on the realizable basis, and compared it to two empirical benchmarks from the Federal Reserve Survey of Consumer Finances 2022 (age cohort 65–74, adjusted for DB pension wealth): the median and the mean actual retirement wealth.

Cohort	Born	Retire	Locked floor	Median actual	Mean actual	vs median
A	1960	2025	$209,942	$260,000	$669,230	0.81×
B	1970	2035	$215,961	$260,000	$669,230	0.83×
C	1980	2045	$229,696	$260,000	$669,230	0.88×
D	1990	2055	$245,435	$260,000	$669,230	0.94×

Table 3. Cohort outcomes, Mode B locked Stable Floor on the realizable basis (central GE return 4.26%, 2025 real dollars). The floor is the sixty-percent locked share; the liquid K3 dividend is reported separately in Section 4.5. Section 6 reports the full distribution from Monte Carlo resampling.

The locked floor lands at approximately the median American's actual retirement wealth — below the median for all four cohorts (0.81×, 0.83×, 0.88×, and 0.94× from earliest to latest). The progression is structural: later cohorts make their deposits against a larger economic base, so the same issuance rules deposit more in real terms. This is not the price-taker reconstruction's several-times-median result; on the realizable basis the floor universalizes the median rather than multiplying it.

4.2 The Median Comparison, Read Correctly

The right way to read a floor that lands at roughly the median is not "it only matches the median." It is that the median saver's outcome becomes the universal minimum. The median is a level half the population never reaches: the Federal Reserve's Survey of Consumer Finances and Distributional

Financial Accounts show the bottom half of households approaching retirement with negligible financial wealth. A floor at the median therefore lifts the entire bottom half of the distribution to a position only the upper-middle currently attains, and it does so for every citizen automatically — with no possibility of undersaving, early withdrawal, fee drag above 0.05 percent, asset concentration, or failure to participate. The framework's contribution is to convert the median from a midpoint that half the population falls below into a floor that none do.

4.3 The Mean Comparison Is the Egalitarian Point

The floor reaches only about 0.31× to 0.37× the mean retirement wealth (Cohorts A through D). Far from a weakness, this is the precise statement of the framework's distributional character. Mean retirement wealth — approximately $669,000 (SCF 2022, age 65–74, DB-adjusted) — is inflated by concentration at the top: the mean is dragged upward by the same wealth concentration the framework is designed to counter. A universal program that delivered the mean would have to reproduce that concentration. The Citizens Standard instead delivers the median to everyone at roughly a third of the mean — the egalitarian result made quantitative. It raises the floor of the distribution without inflating its top, which is what distinguishes a universal capital stake from a mechanism that simply scales existing inequality.

4.4 Sensitivity to the Realizable-Return Band

Because the floor compounds at the general-equilibrium realizable return, the relevant sensitivity is not to price-taker equity scenarios but to the capital-share and depreciation parameters that set that return. The band runs from 3.30 percent (lower capital share, higher depreciation) to 5.03 percent (higher capital share, lower depreciation), with 4.26 percent central. Cohort D, with the longest exposure to the realizable return, illustrates the range:

Scenario	Realizable return	Locked floor	vs median
Low band	3.30% real	$163,606	0.63×
Central	4.26% real	$245,435	0.94×
High band	5.03% real	$345,394	1.33×

Table 4. Cohort D locked floor across the realizable-return band (2025 real dollars). The band reflects general-equilibrium parameter uncertainty (capital share α, depreciation δ) per Neo-Solon, 2026e, not price-taker equity scenarios.

Across the full band the floor moves from roughly three-fifths of the median to about one and a third of it. The same band applied to all four cohorts gives floors of $138K–$298K (A), $143K–$306K (B), $152K–$325K (C), and $164K–$345K (D). Even at the low band the floor reaches roughly half to three-fifths of the median; at the high band it clears the median for every cohort.

4.5 Decomposition, the Liquid Dividend, and Reinvestment

On the realizable basis Cohort A's locked floor decomposes as follows in real (2025) dollars:

Component	Amount (2025$)	Share of floor
K1 deposit at birth (1960)	$816.05	0.39%
K2 floor deposits cumulative (1960–2025)	$39,910.60	19.01%
Total principal (real $)	$40,726.65	19.40%
Compounding gain	$169,215.74	80.60%

Component	Amount (2025$)	Share of floor
Locked Stable Floor	$209,942.40	100.00%

Table 5. Decomposition of Cohort A's locked floor on the realizable basis. The split is roughly 81/19 compounding-to-principal, more principal-weighted than the price-taker reconstruction's 95/5, because the attenuated realizable return does proportionally less of the work.

The shift from a 95/5 to an 81/19 split is itself informative: at the realizable return, compounding does less and the issued principal does more, so the floor depends more directly on the framework's deposits and less on borrowed market performance. The K1 birth endowment, while small in real terms for these old cohorts ($816 to $1,471, or 0.4 to 0.6 percent of the floor, because GDP per capita was low when they were born), is included in every figure; for a citizen born at 2025 launch it is $2,250.

Alongside the locked floor, the citizen receives the liquid K3 dividend — the forty percent of the residual paid as spendable cash. Over a working life this amounts to roughly $26,600 (Cohort A) to $32,500 (Cohort D) in 2025 real dollars. Because it is liquid cash rather than a tax-advantaged locked account, we report it three ways: as cash received; reinvested at an after-tax taxable rate (about 3.4 percent real, reflecting capital-gains and dividend drag on a buy-and-hold position); and at the locked floor's tax-free rate, which is the theoretical ceiling and equals the full residual compounded.

Cohort	Floor (locked)	Dividend as cash	+ cash = total	+ reinvested (after-tax)	+ reinvested (ceiling)
A	$209,942	$26,607	$236,549 (0.91×)	$301,041 (1.16×)	$341,365 (1.31×)
B	$215,961	$27,665	$243,626 (0.94×)	$308,609 (1.19×)	$348,578 (1.34×)
C	$229,696	$29,746	$259,441 (1.00×)	$328,098 (1.26×)	$369,975 (1.42×)
D	$245,435	$32,457	$277,892 (1.07×)	$350,172 (1.35×)	$393,666 (1.51×)

Table 6. The locked floor and the liquid K3 dividend, three ways (multiples are versus median actual wealth). The cash and after-tax-reinvested columns are the load-bearing figures; the tax-free ceiling equals the full residual compounded and is shown only as the maximum. No single behavioral assumption is relied upon: the locked floor stands alone at roughly the median, and the dividend is honestly cash-first with reinvestment as upside.

Read together: the locked floor alone is the guaranteed minimum, at roughly the median. The dividend is current spendable income the median saver scraping by does not have. And a citizen who chooses to reinvest the dividend, even at taxable after-tax rates, clears the median in every cohort — reaching 1.16× to 1.35×. The floor is the promise; the dividend is the option.

4.6 Annual Income at Retirement

We report retirement income only secondarily, because the clean comparison is wealth-to-wealth: the median *income* figure already embeds Social Security, so any income comparison necessarily draws Social Security in on both sides. With that caveat, annual income from a 5 percent draw on the locked floor plus Social Security (full for Cohort A; reduced about 22 percent for Cohorts B–D, who retire after the projected 2032–2034 trust-fund depletion) is:

Cohort	Retire	SS benefit	Median actual income	5% of floor	Floor draw + SS
A	2025	$24,953	$35,353	$10,497	$35,450 (1.00×)
B	2035	$19,463	$29,863	$10,798	$30,261 (1.01×)
C	2045	$19,463	$29,863	$11,485	$30,948 (1.04×)
D	2055	$19,463	$29,863	$12,272	$31,735 (1.06×)

Table 7. Annual retirement income at age 65 (2025 real dollars). Social Security benefits for Cohorts B–D reflect the ~22 percent reduction implied by the projected OASI trust-fund depletion (SSA Trustees, 2026). The floor draw plus Social Security roughly matches median income, but — see Section 4.7 — the floor's durability where Social Security is fragile is the point, not the income parity.

4.7 Social Security as Foil, Not Crutch

The income comparison above is deliberately not the headline, for a structural reason. Social Security, on the 2026 Trustees' projections, faces depletion of its Old-Age and Survivors Insurance trust fund around 2032, with the combined funds depleting by about 2034, after which scheduled benefits would be cut by roughly 22 percent absent congressional action. Its difficulty is demographic: the worker-to-beneficiary ratio has fallen from more than five-to-one in 1960 to 2.9-to-one today and is projected to reach 2.2-to-one by the 2070s. A result that leaned on Social Security to reach median income would inherit that fragility.

The Stable Floor does not. It is funded by seigniorage — new money matched to real growth — rather than by payroll taxes; it is a funded, compounding asset rather than a pay-as-you-go transfer; and it has no worker-to-beneficiary ratio to collapse. It is robust in precisely the dimension where Social Security is fragile. The honest framing is therefore wealth-to-wealth and Social-Security-independent: the floor delivers roughly the median household's capital position to every citizen whether or not Social Security is ever repaired, and it is best read as a hedge against the shortfall rather than a complement that depends on the program's solvency.

4.8 Summary of Findings

The realizable-basis reconstruction supports four deterministic findings, each refined by the Monte Carlo of Section 6:

1. On the realizable basis the locked floor lands at approximately the median actual retirement wealth — 0.81× (Cohort A) rising to 0.94× (Cohort D) — approaching the median rather than multiplying it. Counting the liquid dividend as cash raises the total value captured to 0.91×–1.07× the median.
2. The floor reaches only about 0.31×–0.37× the mean. This is the egalitarian result stated precisely: the framework delivers the median to everyone without reproducing the concentration embedded in the average.
3. Of the locked floor, roughly 81 percent is compounding and 19 percent is deposited principal — more principal-weighted than the price-taker reconstruction, because the attenuated realizable return does less of the work.
4. The result is independent of Social Security and best read as a hedge against its projected shortfall; the liquid dividend, voluntarily reinvested at after-tax rates, lifts every cohort above the median (1.16×–1.35×).

5. Stress Tests and Non-Survivor Analysis

5.1 US Stress Tests: Sequence-of-Returns Scenarios

We applied two adverse historical equity sequences to each cohort, replacing the realizable return during peak accumulation years (ages 25 to 41) with historically observed catastrophic sequences while compounding at the realizable return in all other years. The Great Depression sequence applies Damodaran nominal S&P 500 returns 1929–1945 (deflated to real) to ages 25–41; despite severe nominal losses, cumulative real returns over the window are positive because price-level deflation

between 1929 and 1933 raised the real purchasing power of equity dollars held through the period. The Stagflation sequence applies 1966–1982 returns — seventeen years of flat-to-negative real returns whose cumulative real return is approximately −1 percent, with no compensating deflation. On the realizable basis, where the central path already compounds at 4.26 percent rather than the realized ~6.5 percent, the stress sequences pull the floor below the median for every cohort.

Cohort	Central (realizable)	Depression stress	Stagflation stress	Depr/ cen	Stag/ cen	Depr vs med
A	$209,942	$182,255	$126,470	0.87×	0.60×	0.70× ✗
B	$215,961	$185,966	$131,138	0.86×	0.61×	0.72× ✗
C	$229,696	$195,447	$144,267	0.85×	0.63×	0.75× ✗
D	$245,435	$221,198	$152,428	0.90×	0.62×	0.85× ✗

Table 8. Mode B locked floor under stress scenarios on the realizable basis (2025 real dollars). Depression sequence: 1929–1945 substituted for ages 25–41. Stagflation sequence: 1966–1982 substituted for ages 25–41. Realizable return 4.26% in all non-stress years. Median benchmark: SCF 2022 age 65–74 adjusted for DB pension wealth.

The honest finding is that, on the realizable basis, the hand-picked adverse sequences push the floor below the median for every cohort. Under Stagflation stress — the more punishing of the two, because it had no compensating deflation — the floor falls to 0.49× to 0.59× the median across all four cohorts. Under Depression stress the deflationary real-return boost leaves the floor higher (0.70× to 0.85× the median), though every cohort still finishes below its benchmark. This is a sharper vulnerability than the price-taker reconstruction reported, and it follows directly from compounding at the realizable rather than the realized return: when the central path is already at roughly the median, an adverse accumulation-phase sequence moves a material part of the distribution below it. The structural advantages — universal enrollment, locking, fee minimization, full-horizon compounding — cushion the blow but do not, at the realizable return, keep every cohort above the median through a 17-year stagnation.

5.2 Non-Survivor Analysis

Goetzmann and Jorion (1997) found the US had by far the highest uninterrupted real capital appreciation among major markets, while the median real appreciation rate for other countries was substantially lower. Because the realizable return in this paper is a general-equilibrium marginal product rather than a realized market return, the non-survivor analysis is best read as a check on the framework's relative robustness across institutional environments, not as a transplant of US return magnitudes.

Russia and China experienced complete confiscation of equity markets under revolutionary transitions, producing total loss for citizens; Mode B produces approximately zero effective floor wealth under confiscation, and no financial design can prevent this — the framework's response is constitutional protection. Germany, despite hyperinflation (1921–1923) and two world wars, produced roughly 3.6 percent real equity return over the twentieth century (DMS dataset); the framework's locked, no-withdrawal architecture would have prevented panic-selling during the hyperinflation, a meaningful behavioral advantage. Japan's Nikkei peaked in 1989 and did not recover nominally until 2024, yet the structural elimination of behavioral leakage would have kept a locked account ahead of the median actual Japanese saver. The weak survivors — Belgium, Austria, Spain — produced roughly 1.9 to 2.5 percent real over the century; in such environments the framework's absolute outcomes are proportionally lower while its structural advantages persist.

5.3 Synthesis

The framework's distributional advantage — universalizing the median, lifting the bottom half — is a structural property of universal participation, locking, low fees, and no early withdrawal, not a property of any particular return environment. That structural advantage persists in any market that avoids confiscation, though its absolute magnitude scales with the realizable return available in the implementing economy. Policymakers evaluating non-US implementation should calibrate to their own capital-share and return parameters using the general-equilibrium method of Neo-Solon (2026e), not to US realized returns.

6. Monte Carlo Analysis

The stress tests examine two hand-picked historical sequences. The Monte Carlo answers a distributional question the stress tests cannot: across the full range of plausible histories, what fraction produce a floor below the median actual outcome? On the realizable basis this fraction is substantial, and we report it in full.

6.1 Methodology

The single-path mechanics — K1 at birth, the 60 percent floor share of the growth-matched residual each year, all deposits deflated to December 2025 purchasing power before compounding — are taken unchanged from the deterministic reconstruction. The Monte Carlo changes the return process. Rather than compound at a flat realizable return, each path draws an independent sequence of annual (return, CPI, real-GDP-growth) triples from a historical universe, then recenters the return component so that its geometric (compounding) mean equals the realizable return. This preserves the historical volatility and serial-correlation structure — the sequence-of-returns risk that is the framework's primary financial vulnerability — while shifting the average compounding return down from the realized ~6.5 percent to the realizable 4.26 percent. Geometric (rather than arithmetic) recentering keeps the bootstrap median consistent with the deterministic central case, with volatility drag expressed in the lower tail.

Two universes are used: a wider 1929–2025 window (which includes the Depression and the Great Inflation organically) and a narrower 1960–2025 window matching the deterministic period. Two resampling methods are reported: an IID bootstrap and a moving-block bootstrap (five-year blocks) that preserves within-block sequencing. Each draw is a joint (return, CPI, growth) triple from the same historical year, preserving the empirical correlations — high-inflation years tend to be poor real-return years and to produce smaller K2 deposits simultaneously. Ten thousand paths are drawn per cohort per configuration (twenty thousand for the headline figures), with M2, population, and nominal GDP held at their historical and projected trajectories. The recentering on the realizable return is the one substantive change from a conventional retirement Monte Carlo, and it is what makes the exercise consistent with the general-equilibrium logic of the rest of the paper.

6.2 Results

Table M1 reports the percentile distribution of the locked floor at age 65, in 2025 real dollars, under the block bootstrap drawing from the 1929–2025 universe, recentered on the realizable return.

Cohort	P5	P25	P50	P75	P95	P(<median)	P(<mean)
A (1960)	$47K	$121K	$235K	$451K	$1.18M	54.2%	85.7%
B (1970)	$64K	$158K	$313K	$611K	$1.62M	43.1%	77.8%

Cohort	P5	P25	P50	P75	P95	P(<median)	P(<mean)
C (1980)	$80K	$199K	$390K	$787K	$2.17M	34.3%	69.9%
D (1990)	$94K	$231K	$454K	$916K	$2.48M	29.0%	64.9%

Table M1. Monte Carlo locked-floor distribution, 1929–2025 universe, block bootstrap (5-yr blocks), returns recentered on the realizable geometric mean. 10,000 paths per cohort. P(<median) and P(<mean) are the empirical probabilities the floor falls below the SCF 2022 benchmark adjusted for DB pension wealth.

Three findings follow directly. First, the typical (P50) floor runs from $235K (Cohort A) to $454K (Cohort D), or 0.90× to 1.75× the median — the later cohorts' longer compounding windows push their median outcomes well above the benchmark, the earliest cohort's lands just below it. Second, the share of paths below the median is large and declining across cohorts: 54 percent (A), 43 percent (B), 34 percent (C), 29 percent (D). This is the honest distributional cost of the realizable return, and it is far higher than the 1–7 percent a price-taker reconstruction would report. Third, the share below the mean is higher still (65–86 percent), which is the egalitarian point restated: the floor is built to deliver the median, not the mean, so most paths sit below the mean by construction.

These figures are for the locked floor alone. Counting the liquid dividend voluntarily reinvested at an after-tax rate shifts the whole distribution up: the share of paths below the median falls to roughly 35 percent (A), 21 percent (B), 11 percent (C), and 6 percent (D), and the P50 clears the median for every cohort. The locked floor is the guaranteed distribution; the floor-plus-reinvested-dividend is the upper distribution; the truth for any given citizen lies between, depending on how much of the dividend they spend.

6.3 Configuration Sensitivity

Table M2 confirms the findings do not depend on universe or method.

Cohort	Configuration	P50	P5	P(<median)
A	1929–2025, IID	$234K	$51K	54.0%
A	1929–2025, Block	$235K	$47K	54.2%
A	1960–2025, IID	$235K	$58K	54.4%
A	1960–2025, Block	$223K	$62K	57.9%
B	1929–2025, Block	$313K	$64K	43.1%
B	1960–2025, Block	$282K	$74K	46.3%
C	1929–2025, Block	$390K	$80K	34.3%
C	1960–2025, Block	$341K	$87K	37.6%
D	1929–2025, Block	$454K	$94K	29.0%
D	1960–2025, Block	$384K	$93K	33.3%

Table M2. Sensitivity of the locked floor (P50, P5, and P(<median)) to bootstrap configuration. 10,000 paths per row, returns recentered on the realizable mean. The qualitative findings are stable across universe and method.

The narrower 1960–2025 universe slightly tightens the tails and modestly raises the below-median share for the earliest cohort; block versus IID makes little difference to the median or the below-median probability. A finding that requires explicit acknowledgment: in every configuration the P5 floor (roughly $47K to $94K depending on cohort) is well below the median benchmark. Across the worst paths the floor delivers a fraction of the median, which is a property of any equity-based vehicle exposed to genuine sequence risk. The framework's defense against this tail is structural — the architectural

paper's Equity Market Stability Reserve (Tool 12) and the diversification of total-market exposure — and the bootstrap is run without those defenses activated, so the reported tail is conservative.

6.4 What the Monte Carlo Establishes

The bootstrap establishes four things on the realizable basis. First, the locked floor's typical outcome rises from just below the median for the earliest cohort to roughly 1.75 times the median for the latest, as longer compounding windows reduce sequence risk — Cohort D, the most-projected cohort, has the most reliable median-beating record, inverting a casual reading of the deterministic results. Second, the floor falls below the median in a substantial minority-to-majority of paths (29 to 54 percent across cohorts), so the framework's distributional claim is best stated as *in the central tendency and for later cohorts*, not unconditionally. Third, counting the liquid dividend reinvested at after-tax rates lifts every cohort's median above the benchmark and cuts the below-median share to 6–35 percent. Fourth, the adverse tail is real and structural rather than statistical; no equity-based vehicle eliminates it. The Monte Carlo does not establish superiority over alternative reforms, does not address transition or implementation, and does not characterize the withdrawal phase. What it establishes is that the realizable-basis floor universalizes roughly the median in the central case, with an honestly large dispersion that the framework's structural defenses, omitted here, are designed to bound.

6.5 Replication Note

The Monte Carlo engine extends the same K1/K2 deposit formula used in the deterministic reconstruction, with the return component recentered on the realizable geometric mean. The engine, historical data files (including the 1929–1959 extension), and figure-generation scripts are released as authoritative_data.py, deterministic_engine.py, mc_engine.py, and run_ge_results.py. Random seeds are fixed; results are reproducible on Python 3.8+ with numpy. The realizable return and the 60/40 floor share are exposed as the parameters GE_REALIZABLE_RETURN and FLOOR_SHARE, so a reader can reproduce the price-taker reconstruction (by disabling the recentering) or vary the realizable return across its band.

7. The Result in Context: Capital Ownership, Not Retirement

This paper measures the Stable Floor against the capital that Americans actually accumulate over a lifetime. That comparison draws on retirement-balance data because it is the best available record of end-of-life household wealth — but the object being evaluated is not a retirement product. The Stable Floor is a permanent, individually owned equity stake that becomes liquid at sixty-five and remains heritable beyond it; the relevant literature is therefore the distribution of capital ownership and the effects of asset-holding, not the adequacy of retirement saving. The architectural paper (Neo-Solon, 2026a) situates the framework among monetary-reform alternatives; this section situates the empirical result among the wealth-distribution and asset-based-welfare literatures it is benchmarked against.

The motivating fact is the concentration of capital ownership. The Federal Reserve's Distributional Financial Accounts and Survey of Consumer Finances show that corporate equity and mutual-fund wealth is held overwhelmingly by the top decile of households, with the bottom half owning a negligible share; Wolff (2017) and Saez and Zucman (2016) document the long-run trajectory of this concentration. This is exactly why the median-versus-mean distinction of Section 4.3 is the heart of the result: the mean retirement balance is an artifact of that concentration, and a floor at the median is a floor that deliberately does not reproduce it. The wealth gap is, more precisely, an asset gap rather than an income gap: Oliver and Shapiro (1995) showed that disparities in accumulated assets dwarf and outlast disparities in income, and Hamilton and Darity (2010) framed the racial wealth gap as a

problem of capital endowment that income transfers cannot close. A universal equity stake speaks to this literature in a way that an income transfer does not, which is the empirical significance of the result reported here.

The Stable Floor's nearest conceptual relatives are the stakeholding and asset-based-welfare proposals. Ackerman and Alstott (1999) proposed a universal capital grant to every young citizen; Atkinson (2015) proposed a minimum inheritance, or universal capital endowment, paid at adulthood; and Piketty (2020) proposed an inheritance for all financed by progressive wealth and inheritance taxation. Sherraden (1991), in the foundational statement of asset-based welfare, argued that owning an asset changes a household's security, time horizon, and behavior in ways an income-equivalent transfer does not — the precise distinction between the locked floor and the liquid dividend, and the reason the framework treats the locked equity stake as primary and the cash dividend as secondary. The political-theory lineage runs through the property-owning-democracy tradition of Meade (1964) and Rawls (2001), which holds that distributing productive capital broadly is categorically different from redistributing income after the fact.

What distinguishes the Stable Floor within this literature is its funding source and its form. The stakeholding proposals finance the endowment through taxation; the Stable Floor is funded by monetary seigniorage routed to citizens at issuance, requiring no tax and adding no government liability — the same property that makes it immune to the demographic arithmetic now threatening Social Security. And where most proposals grant a one-time sum or a collectively managed fund, the Stable Floor is an individually owned account that compounds across a lifetime under a constitutional lock. The empirical question this paper answers is therefore narrow and concrete: given the actual distribution of capital ownership and the actual lifetime accumulation of median households, how large a stake would the framework's issuance rules have produced on a realizable-return basis, and how does it compare to what households hold today. The answer — roughly the median, for everyone, at a fraction of the mean — is a statement about ownership, with the retirement figures serving only as the yardstick.

8. The Transition Period: A Forward Projection

Sections 3 through 6 reconstruct what the mature system would have delivered over the actual 1960–2025 record on the realizable basis. This section addresses a forward-looking question: what would a citizen born at or after enactment receive, living through the debt-retirement transition (Neo-Solon, 2026c) and into the permanent Mode T-stable steady state? Because no citizen has yet lived this path, the results here are projections, labeled as such and never blended with Part I.

Why this section is a companion to the transition paper. The transition paper specifies how pre-existing public debt is retired: a Legacy Debt Trust refinances the debt held by the public while a transition-only issuance channel, KT, directs price-level-calibrated money creation to bond redemption. The KT channel never deposits into citizen accounts — it redeems bonds held by financial institutions. A citizen living through the transition therefore accumulates K1 and the 60 percent floor share exactly as in the mature steady state. The only difference is second-order: during the debt-paydown window the systematic equity demand created by the architecture compresses forward returns by an estimated 0.4 to 0.6 percentage points (transition paper, Section 4.3). This section quantifies that effect on the realizable basis.

8.1 Method

We model four forward cohorts born in 2026, 2036, 2046, and 2056, each accumulating from birth to age 65 (retirement years 2091, 2101, 2111, 2121). Each receives K1 at birth and the 60 percent floor

share annually. Two differences distinguish this from Part I. First, all returns are the realizable return rather than historical, reported across the realizable band — low (3.30 percent), central (4.26 percent), and high (5.03 percent). Second, during the debt-paydown window (2026 through approximately 2071) a 0.5 percentage point compression is applied, reverting to the unadjusted realizable return once the window closes. These are projections conditional on the stated assumptions.

8.2 Forward Cohort Outcomes

Cohort	Born	Retire	Low band (3.30%)	Central (4.26%)	High band (5.03%)
T1	2026	2091	$222,624	$327,986	$455,997
T2	2036	2101	$268,606	$396,372	$551,561
T3	2046	2111	$322,106	$476,420	$663,935
T4	2056	2121	$383,531	$568,696	$793,933

Table 9. Projected locked floor at age 65 for forward transition cohorts on the realizable basis, with a 0.5 percentage point compression during the 2026–2071 paydown window. All figures in 2025 real dollars. Forward cohorts accumulate against a larger future economic base, so their floors exceed the historical cohorts'; the liquid dividend is additional. These are forward projections, not historical reconstructions.

8.3 The Cost of the Transition to Citizens

Because the KT channel retires debt without drawing on citizen accounts, the only cost a citizen bears is the modest compression during the paydown window. Table 10 isolates it by comparing each cohort's central-case floor with and without compression.

Cohort	Born	Without compression	With compression	Transition cost
T1	2026	$370,511	$327,986	−11.5%
T2	2036	$430,599	$396,372	−7.9%
T3	2046	$500,470	$476,420	−4.8%
T4	2056	$581,725	$568,696	−2.2%

Table 10. The cost of the transition to citizens, central case on the realizable basis. The compression is largest for the earliest cohort, whose entire early-accumulation period falls within the paydown window, and declines monotonically as later cohorts experience proportionally less of their accumulation under compression.

Interpretation. Even for the most exposed cohort the transition reduces the projected floor by roughly one-ninth; for cohorts born thirty years in, the cost falls below one-fortieth. This is the empirical counterpart to the transition paper's structural claim: the debt is retired through the KT channel and the primary surplus, not through any reduction in citizen accumulation. The transition does not ask citizens to fund the paydown out of their floors; it asks only that they accept a temporary, declining return compression that disappears once the debt is retired and the system settles into Mode T-stable.

Honesty boundary. These projections rest on assumptions history may not honor. The 4.26 percent central realizable return is a general-equilibrium estimate, not a guarantee; the 0.5 percentage point compression is a modeled estimate whose precise magnitude is an open question (transition paper reports 0.4 to 0.6 percentage points). The low-band case, which combines a 3.30 percent realizable return with the compression, is the appropriate figure for a cautious reader. What the projection establishes is not a dollar amount but a structural conclusion robust across the band: the transition's

cost to citizens is modest and declining, because the debt-retirement mechanism does not draw on citizen accounts.

9. Discussion

9.1 What the Data Establishes

On the realizable basis the framework places every citizen at approximately the median actual retirement wealth — just below it for every cohort on the locked floor alone — and well below the mean. That is a distributional claim, not a claim to outsized returns: it lifts the bottom half of the wealth distribution to a position only the upper-middle currently reaches, universally and automatically, while deliberately not reproducing the concentration that inflates the mean. The Monte Carlo refines this: the typical (P50) floor runs from 0.90× the median (Cohort A) to 1.75× (Cohort D), with a substantial minority-to-majority of paths (29 to 54 percent) falling below the median, reduced to 6–35 percent when the liquid dividend is reinvested. The framework's strongest defensible claim is therefore distributional and conditional: it universalizes roughly the median in the central case and for later cohorts, with honest dispersion around that center.

The decomposition shows that roughly 81 percent of the floor is compounding and 19 percent is deposited principal — a more principal-weighted split than a price-taker reconstruction implies. The framework does not invent the compounding; it guarantees the conditions — universal participation, automatic deposits, locked investment, fee minimization, no early withdrawal — under which compounding can occur for every citizen regardless of discipline or sophistication. Whether an equivalently structured but differently funded vehicle would reproduce the result is the open question stated next.

9.2 What the Data Does Not Establish

The analysis does not establish that the Citizens Standard is superior to all alternative reforms; the comparator is median and mean actual outcomes under the current US system. It does not resolve whether any automatic-deposit, locked-investment vehicle funded differently would produce similar results. The general-equilibrium and price-level consequences of the framework — including the realizable-return result this paper imports rather than derives — are formalized separately in the macroeconomic model (Neo-Solon, 2026e). It does not address transition costs or implementation feasibility beyond Part II, which are treated in the transition (Neo-Solon, 2026c) and statutory implementation (Neo-Solon, 2026d) papers. The Monte Carlo characterizes accumulation-phase uncertainty but not withdrawal-phase outcomes, which depend on post-65 sequence risk not modeled here. Finally, the reconstruction measures floor accumulation from the K1 and K2 ownership channels; the liquid K3 dividend is reported as cash and as an optional reinvestment, but the inflation-gap stabilizer KI and other current-income channels distribute spendable income rather than locked floor wealth and fall outside the cohort-accumulation scope. How those channels behave on prices — whether the framework's issuance rule and the KI self-correction actually contain an inflation rather than merely target one — is taken up empirically in the inflation counterfactual reported in Appendix A.9, which runs Tool 14 against the 2021–2023 surge and the 1972–1983 Great Inflation and draws its transmission from the macroeconomic model's Proposition 6 (Neo-Solon, 2026e). It is a companion to the cohort result, not a component of it.

9.3 Relationship to the Architectural Paper's Claims

The architectural paper's claim that Mode B reliably produces retirement wealth in excess of median American outcomes must be restated on the realizable basis. What the data supports is narrower and

more defensible: Mode B universalizes roughly the median (0.81× to 0.94× on the locked floor alone), reaching or exceeding it only with the liquid dividend or in the central tendency of later cohorts, with a non-trivial share of adverse paths falling short — a claim about distribution rather than multiples. The architectural paper's published Mode B figures are forward projections from 2025 launch parameters and are larger in absolute magnitude because of the larger contemporary base and, where they used a price-taker return, because of the return assumption this paper corrects. Both sets of figures are coherent on their own terms once the return basis is made explicit; this paper's contribution is to make it explicit and to report the realizable result without inflation.

9.4 Honest Positioning

The Citizens Standard, evaluated empirically through its age-65 balances on the realizable basis, presents as a structural floor on lifetime capital ownership at roughly the median household's level — universal, self-funded, demographically immune, and constitutionally locked — with meaningful dispersion and genuine below-median tail risk that no equity-based vehicle escapes. Its principal social contribution is distributional: eliminating the savings-discipline lottery and the participation gap so that the median outcome becomes the floor, not improving the absolute returns available to disciplined high-end savers and not eliminating equity-market tail risk. Stated as multiples it is modest; stated as a distribution it is radical.

10. Conclusion

This paper set out to evaluate the Citizens Standard empirically — to ask not whether the framework is theoretically coherent but what outcomes it would have produced against historical data, once its returns are evaluated on a basis a universal program can actually achieve. The answer is more modest in magnitude and more radical in distribution than a naïve reconstruction suggests.

Because the framework is universal, it cannot earn the price-taker return that atomistic investors realized on a capital stock their saving did not move. Compounding every cohort's deposits at the general-equilibrium realizable return of 4.26 percent (band 3.30 to 5.03 percent) rather than the realized ~6.5 percent, the locked Stable Floor lands at approximately the median American's actual retirement wealth — 0.81× for the earliest cohort rising to 0.94× for the latest — and at only about 0.31× to 0.37× the mean. That is the result, stated honestly: not several times the median, but the median universalized, at a fraction of a mean that is itself an artifact of concentration.

The distributional reading is what makes this significant. The median is a level half the population never reaches; a universal floor at the median lifts the entire bottom half to a position only the upper-middle attains today, automatically and without the savings-discipline lottery that otherwise determines whether a citizen reaches sixty-five with meaningful capital. The liquid K3 dividend adds spendable income the median saver does not have, and, if voluntarily reinvested even at taxable after-tax rates, lifts every cohort above the median. Roughly 81 percent of the floor is compounding and 19 percent is the issued principal; the framework guarantees the conditions under which the compounding occurs for everyone.

The framework does not perform uniformly. Under stagflation-magnitude sequences during peak accumulation the floor falls to 0.49× to 0.59× the median; under Monte Carlo resampling 29 to 54 percent of paths fall below the median, and the worst paths deliver a fraction of it. These are honest costs of evaluating returns realistically, and the framework's defenses against them are structural rather than statistical. Crucially, the result is independent of Social Security — which on the 2026 Trustees' projections faces a roughly 22 percent benefit cut around 2032–2034 — because the floor is

funded by seigniorage rather than payroll taxes and carries no demographic fragility. It is best read as a hedge against that shortfall.

The Citizens Standard is a serious policy proposal that deserves serious empirical scrutiny. On the realizable basis the data supports a real and defensible contribution: a structural floor on lifetime capital ownership at roughly the median household's level, self-funded, demographically immune, and universal, whose effectiveness derives from eliminating the behavioral and institutional failures that otherwise determine whether a citizen reaches retirement with meaningful capital. Modest in multiples, radical in distribution — that is the claim the framework can honestly make, and it is enough.

References

Ackerman, Bruce, and Anne Alstott. The Stakeholder Society. Yale University Press, 1999.

Atkinson, Anthony B. Inequality: What Can Be Done? Harvard University Press, 2015.

Bengen, William P. “Determining Withdrawal Rates Using Historical Data.” Journal of Financial Planning, October 1994.

Braggion, Fabio, Rik Frehen, and Michael Haliassos. “Inflation and Individual Investors’ Behavior: Evidence from the German Hyperinflation.” Review of Financial Studies, 2022.

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.

Committee for a Responsible Federal Budget. Analysis of the 2026 Social Security Trustees’ Report. 2026.

Congressional Budget Office. The Long-Term Budget Outlook. 2025.

Damodaran, Aswath. “Historical Returns on Stocks, Bonds and Bills: 1928–2025.” NYU Stern, January 2026. Available at pages.stern.nyu.edu/~adamodar.

Dimson, Elroy, Paul Marsh, and Mike Staunton. Triumph of the Optimists: 101 Years of Global Investment Returns. Princeton University Press, 2002.

Efron, Bradley, and Robert J. Tibshirani. An Introduction to the Bootstrap. Chapman & Hall, 1993.

Federal Reserve Bank of St. Louis. M2 Money Supply (M2SL). FRED, 2025.

Federal Reserve Board. Distributional Financial Accounts. 2025.

Federal Reserve Board. Survey of Consumer Finances, 2022.

Goetzmann, William N., and Philippe Jorion. “A Century of Global Stock Markets.” NBER Working Paper No. 5901, 1997.

Hamilton, Darrick, and William Darity. “Can ‘Baby Bonds’ Eliminate the Racial Wealth Gap?” Review of Black Political Economy, Vol. 37, 2010.

Künsch, Hans R. “The Jackknife and the Bootstrap for General Stationary Observations.” Annals of Statistics, Vol. 17, No. 3, 1989.

Meade, James E. Efficiency, Equality and the Ownership of Property. Allen & Unwin, 1964.

Neo-Solon (2026a). The Citizens Standard: One Model, Many Systems — A Constitutional Monetary Architecture. SSRN Working Paper 6702518.

Neo-Solon (2026c). The Citizens Standard: Transition Architecture and Migration Mechanics. SSRN Working Paper 6810741.

Neo-Solon (2026d). The Citizens Standard: A Statutory Implementation Pathway. SSRN Working Paper 6873798.

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

Oliver, Melvin, and Thomas Shapiro. Black Wealth/White Wealth. Routledge, 1995.

Piketty, Thomas. Capital and Ideology. Harvard University Press, 2020.

Rawls, John. Justice as Fairness: A Restatement. Harvard University Press, 2001.

Saez, Emmanuel, and Gabriel Zucman. “Wealth Inequality in the United States since 1913.” Quarterly Journal of Economics, Vol. 131, No. 2, 2016.

Sherraden, Michael. Assets and the Poor: A New American Welfare Policy. M.E. Sharpe, 1991.
Shiller, Robert J. Online Data. Yale University, 2025. Available at econ.yale.edu/~shiller/data.htm.
Siegel, Jeremy J. Stocks for the Long Run, 5th edition. McGraw-Hill, 2014.
Social Security Administration. The 2026 Annual Report of the Board of Trustees. June 2026.
US Census Bureau. National Population Estimates, Vintage 2025 (January 27, 2026 release).
Vanguard Group. How America Saves 2025. Vanguard Research, 2025.
Wolff, Edward N. A Century of Wealth in America. Harvard University Press, 2017.

Technical Appendix

The Citizens Standard: A Historical Counterfactual — Empirical Analysis of an Alternative US Monetary Architecture, 1960–2055. Neo-Solon (2026) · Working Paper · SSRN 6735078.

This appendix provides the complete mathematical specification of every quantitative result, on the realizable-return basis used throughout the paper. All results are reproducible from the Python replication package released as a supplementary file on SSRN and archived on GitHub (Neo-Solon/Citizens-Standard).

A.1 Notation and Variable Definitions

Symbol	Definition / source
t	Calendar year index
b(i), r(i)	Birth year and retirement year (= b(i)+65) of citizen i
GDP(t), N(t)	Nominal US GDP (BEA/FRED GDPA); resident population (Census)
g(t)	Real GDP growth rate (BEA A191RL1A225NBEA)
M2(t)	M2 money supply, end-of-period December (FRED M2SL)
CPI(t), π(t)	CPI-U Dec-Dec index and inflation rate (BLS)
r*	Realizable real return on the floor (Neo-Solon 2026e, §6.7): 4.26% central, band 3.30–5.03%
φ	Floor share of the post-K1 residual (Mode B 60/40): φ = 0.60
K1(t), K2(t)	K1 deposit per new citizen; K2 floor deposit per citizen (formulas A.2)
K3(t)	Liquid citizen dividend per citizen = (1−φ)/φ × K2(t)
SF(i)	Locked Stable Floor at retirement, 2025 real USD (formula A.3)

A.2 The K Issuance Formulas

The framework issues new money through two ownership channels. K1 is triggered by verified citizenship events; the annual growth-matched line is triggered by positive real growth and, under Mode B, split 60/40 into the locked floor (K2) and the liquid dividend (K3).

A.2.1 K1 — Citizenship Endowment

K1(t) = α · gdppc(t) = 0.025 · GDP(t)/N(t), α = 0.025 (Mode B).

At launch: K1(2025) = 0.025 × $90,000 = $2,250 per new citizen. A citizen naturalizing at age a receives (65−a)/65 of the full endowment.

A.2.2 K2 and K3 — Growth-Matched Line, Split 60/40

residual(t) = β · g(t) · M2(t−1)/N(t) − K1 deduction, β = 1.0, if g(t) > 0; else 0.

K2(t) = φ · residual(t) = 0.60 · residual(t) (locked floor)

K3(t) = (1 − φ) · residual(t) = 0.40 · residual(t) (liquid dividend, not compounded into the floor)

The full growth-matched line (β = 1.0) holds the price level stable by construction (money grows with real output); the 60/40 split allocates that line between locked floor wealth and current spendable income. Only K1 and K2 enter the cohort-accumulation model below; K3 is reported separately in Section 4.5.

A.3 The Stable Floor — Cohort Accumulation Model

Real-return basis. Deposits are deflated to December 2025 dollars before compounding, and compounding uses the realizable real return r* (not the realized series):

$D_{real}(t) = D_{nominal}(t) \cdot CPI(2025)/CPI(t)$

$Balance(t+1) = [Balance(t) + K1_{real}(t) + K2_{real}(t)] \cdot (1 + r^*)$

with K1 deposited only in the birth year and K2 each year real growth is positive. Deposit timing is beginning-of-period (deposit-then-compound): each year's deposits are credited at the start of the year and earn that year's return. In the Monte Carlo of Section 6, the constant r* is replaced by a bootstrapped sequence whose geometric mean equals r* (Section A.5).

A.4 The Decomposition Algebra

Principal and compounding shares are defined as:

$P_{real}(i) = K1_{real}(b(i)) + \Sigma_t K2_{real}(t) \cdot 1[g(t)>0]; \quad Compounding\ share = 1 - P_{real}(i)/SF(i).$

Cohort A (born 1960), realizable basis:

Component	Amount (2025$)	Share
K1 deposit at birth (1960)	$816.05	0.39%
K2 floor deposits cumulative (1960–2025)	$39,910.60	19.01%
Total real principal	$40,726.65	19.40%
Compounding gain	$169,215.74	80.60%
Locked Stable Floor	$209,942.40	100.00%

The 81/19 split (versus 95/5 under a price-taker return) is a direct consequence of the lower realizable return: at 4.26 percent rather than ~6.5 percent, compounding does proportionally less of the work and the issued principal does more.

A.5 Bootstrap Monte Carlo Methodology

The bootstrap draws joint triples $\omega(t) = (R(t), \pi(t), g(t))$ from a historical universe (1929–2025 or 1960–2025), preserving the empirical correlations among returns, inflation, and growth. Two methods are reported: IID and moving-block (5-year blocks, circular wraparound).

Realizable recentering. The drawn real-return sequence is rescaled so its geometric mean equals r*. If g_{uni} is the universe geometric mean real return, each drawn gross return $(1+R)$ is multiplied by $(1+r^*)/(1+g_{uni})$. This preserves volatility and serial correlation while shifting the average compounding return to the realizable level. Geometric (rather than arithmetic) recentering keeps the bootstrap median aligned with the deterministic central case, with volatility drag appearing in the lower tail. The K2 zero-floor ($K2 = 0$ when $g \leq 0$) is applied per drawn year. 10,000–20,000 paths per configuration; numpy default_rng with fixed seed; results stable across seeds.

A.6 Full Parameter Table

Parameter	Value	Source
K1 fraction (Mode B)	$\alpha = 0.025$	Architectural paper, Mode B
Floor share of residual	$\varphi = 0.60$ (60/40 split)	Architectural paper, Mode B

Parameter	Value	Source
Growth-match coefficient	β = 1.0	Full real-growth-matched line
Realizable real return (central)	r* = 4.26%	Neo-Solon 2026e, §6.7 (GE fixed point)
Realizable return band	3.30% – 5.03%	Capital-share / depreciation range
K3 reinvestment (after-tax)	3.4% real	Illustrative taxable buy-and-hold drag
K1 deposit at launch	$2,250 / new citizen	0.025 × $90,000 gdppc
K2 floor zero-floor	K2 = 0 if g(t) ≤ 0	Structural design rule
Nominal GDP (2025)	$30,762B	BEA GDPA Apr 2026
Population (2025)	341.8M	Census Vintage 2025
M2 end-of-period Dec (2025)	$22,366B	FRED M2SL
Real GDP growth (2025)	2.1%	BEA A191RL1A225NBEA
Post-2025 nominal GDP / M2 growth	4.0% / 4.5%	CBO LT Outlook 2025
Post-2025 population / real GDP / CPI	0.4% / 1.8% / 2.5%	CBO LT Outlook 2025
Paydown-window compression	0.5 pp, 2026–2071	Transition paper §4.3 (0.4–0.6 pp)
Bootstrap universe / block size	1929–2025 / 5 yrs	Damodaran / BLS / BEA
Paths per configuration	10,000–20,000	Convergence verified
Retirement age	65	SSA Full Retirement Age

A.7 Stagflation Stress Sequence

The 1966–1982 block is applied as the stagflation stress (Section 5.1), substituted for ages 25–41. Its defining property is a cumulative 17-year real return of approximately −1 percent — essentially zero real growth across the full window, with no compensating deflation. Four years meet both high-inflation (CPI > 5%) and low-growth (real GDP < 2%) conditions simultaneously: 1970, 1974, 1975, and 1980. The Depression stress applies 1929–1945; despite severe nominal losses its cumulative real return is positive because of 1929–1933 deflation. In both stress tests, all non-stress years compound at the realizable return r*.

A.8 Replication Code Reference

All results are reproduced by the Python replication package (GitHub: Neo-Solon/Citizens-Standard).

Script	Contents	Reproduces
authoritative_data.py	Source-of-truth historical data, 1928–2025	All historical inputs
deterministic_engine.py	K1/K2 formulas; 60/40 split; realizable-return compounding	Tables 1–10, Sections 3–5
mc_engine.py	Block bootstrap with realizable recentering	Section 6 Monte Carlo
run_ge_results.py	Runs every table on the realizable basis	Full paper replication
transition_cohorts.py	Forward cohorts with paydown compression	Tables 9–10

Key parameters exposed: GE_REALIZABLE_RETURN (default 0.0426; set across the 0.0330–0.0503 band) and FLOOR_SHARE (default 0.60). Setting GE_REALIZABLE_RETURN to None and disabling the Monte Carlo recentering reproduces the price-taker reconstruction, so a reviewer can

audit the difference the realizable basis makes. Runtime for the full grid is approximately one second; requirements are Python 3.8+, numpy, matplotlib.

The inflation counterfactual is packaged in the same repository as a self-contained module (tool14_inflation) with its own scripts and data; it is described in Appendix A.9.

A.9 Inflation Counterfactual (Tool 14)

This paper measures the framework's effect on cohort wealth; its effect on the price level is a separate question, addressed here as a companion. The inflation counterfactual runs the framework's Tool 14 against two documented US inflations: the 2021–2023 surge (BLS CPI-U, with the demand-versus-supply split taken month-by-month from the San Francisco Fed's published decomposition) and the full 1972–1983 Great Inflation (computed from the BLS monthly index, so the framework is present through the build-up rather than dropped in at the peak). Two counterfactuals are kept distinct: prevention, in which the framework is the system throughout, and response, a drop-in once inflation is already high. The headline result is reported two independent ways — a transparent reduced form that removes the demand-driven share and applies Tool 14 at its capacity, and the macroeconomic model's own Proposition 6 transmission (Neo-Solon, 2026e), in which KI self-corrects the price gap. The two agree to within roughly one percentage point: prevention holds the framework's inflation near 4 percent in both episodes, against actual peaks of 9.1 percent (June 2022) and 14.8 percent (March 1980). The exercise does not claim the framework prevents a supply shock; it claims a lower, self-correcting peak reached without an interest-rate channel, and it is explicit that Tool 14's disinflation is slower than an aggressive rate shock, not faster. Scripts run_counterfactual.py (reduced form), structural_run.py (Proposition 6 transmission), and sensitivity.py reproduce the figures.