

THE CITIZENS STANDARD

Distribution and Wealth Inequality

A Microsimulation Grounded in the 2022 Survey of Consumer Finances — How a Universal Locked Floor Reshapes the Distribution from Below, and Where It Does Not

Neo-Solon

Neo-Solon@hotmail.com

WORKING PAPER — JUNE 2026

COMPANION PAPERS

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

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

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

The Citizens Standard: The Structural Buyer (Neo-Solon, 2026h) · SSRN 6945320

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

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

Abstract

Inequality is the problem the Citizens Standard (CS) is built to address, yet distributional claims are the easiest to overstate. This paper measures the effect on the wealth distribution honestly, using a microsimulation seeded from the actual 2022 Survey of Consumer Finances (SCF) household records rather than a synthetic distribution. Four channels are modelled, each tied to data or to a named causal study: the locked floor (the engine's deterministic per-citizen accumulation), the standing dividend, the structural buyer's return compression, and bequests. The central finding is bounded and deliberately not utopian. The floor transforms the bottom of the distribution: it abolishes the zero-and-negative-net-worth tail (from 7.9% of households to about 1.5%), lifts the 10th percentile from roughly \$450 to roughly \$65,000, and nearly triples the share of all wealth held by the bottom half (2.2% to 6.7%). The wealth Gini falls from 0.83 to about 0.74, but the decomposition is unambiguous: roughly three-quarters of that reduction is the floor lifting the bottom, and the top-1% share moves only a few points, mostly mechanically. Return compression — grounded in the measured returns-to-wealth gradient (Fagereng et al. 2020) and the inelastic-markets multiplier (Gabaix & Koijen 2021), bounded by the structural-buyer paper — trims the top-1% share by a further one to four points over a generation, and is the most model-dependent channel. The bequest result is the most striking: registry studies (Nekoei & Seim 2023; Elinder et al. 2018) find that ordinary inheritances equalise on receipt but the effect reverses within a decade because poorer heirs deplete the windfall and wealthier heirs earn higher returns on it; the CS floor bequest is structurally immune to both mechanisms because it is locked and earns a uniform return, so it converts inheritance's short-lived equalising effect into a permanent one. The honest summary: CS is a powerful floor-raiser with bounded top-compression. It rebuilds the floor of the distribution and universalises inheritance; it does not level the top, and a 0.74 Gini remains a high-inequality society. Every result is consistent with the bounded-ownership ceiling of the structural-buyer paper and survives the obvious robustness checks.

Methods. The baseline is the weighted 2022 SCF (4,595 families, five implicates), verified to reproduce the published mean (\$1.06M), median (\$192,700), top-1% threshold, and a 0.830 wealth Gini before any CS channel is applied. The floor is the issuance-engine's deterministic accumulation by age, deflated to 2022 dollars; the other channels are anchored as described above. Results are staged by robustness — mechanical first, model-dependent last — and reported with a sensitivity band.

JEL classification: D31, D63, E21, H23, G51

Keywords: Citizens Standard; wealth inequality; wealth distribution; Gini; microsimulation; Survey of Consumer Finances; universal floor; inheritance; returns to wealth

Contents

1. The Question, and Why It Is Hard.....	4
2. What the Series Already Implies — and What Is Missing.....	4
3. The Model.....	4
4. Data and Calibration.....	5
5. Results.....	6
5.1 The Floor: the Robust Result.....	6
5.2 The Dividend.....	6
5.3 Return Compression: Bounded and Secondary.....	6
5.4 Bequests: Near-Universal and Permanently Equalising.....	7
5.5 Decomposition: Floor versus Top.....	7
6. Robustness.....	7
7. Honest Limits.....	8
8. Distributional Claims.....	8
Technical Appendix.....	10
A. The Microsimulation.....	10
B. Consistency with the Structural Buyer.....	10
C. Sensitivity Table.....	10
References.....	11

1. The Question, and Why It Is Hard

The motivating problem, and the temptation that comes with it. Concentrated wealth is the condition the Citizens Standard is meant to change, so the natural question is how much it changes it. The temptation is to answer triumphantly — to claim the architecture dissolves the wealth gap. This paper resists that, for a reason internal to the series: the structural-buyer paper (Neo-Solon 2026h) already proved that universal ownership is bounded, capped at roughly six to twelve per cent of the market by the c/g ceiling. That six-to-twelve-per-cent figure is the US calibration, and the load-bearing bound is not the point estimate but the c/g ceiling itself, which rises with a shallower equity market or a faster-growing economy; in a case like India the realized share is closer to half. The distributional analysis in this paper is US-specific throughout, benchmarked to US wealth data, so the US bound is the relevant one here, but the ownership share is a property of the implementing economy rather than a universal constant. A paper that claimed CS crushes the top would contradict our own result. The honest question is therefore narrower and more answerable: through which channel, and by how much, does a universal locked floor reshape the distribution — and where does it leave the top untouched?

Why distributional claims are hard to make honestly. Inequality results are notoriously sensitive to modelling choices — the income process, saving behaviour, return heterogeneity, bequest rules. A model with enough free parameters can produce almost any Gini path. The discipline this paper adopts is to seed from real household records, to tie every channel to data or to a named causal study, to stage results by how mechanical they are, and to report the model-dependent channels with explicit bands. Where a result rests on an assumption, the assumption is stated and varied.

What this paper claims, in one sentence. CS is a powerful floor-raiser with bounded top-compression: it abolishes the bottom tail, lifts the lower half of the distribution, and universalises inheritance, while the top moves only modestly — so the Gini falls a real but undramatic amount, driven overwhelmingly from below.

2. What the Series Already Implies — and What Is Missing

Three papers already speak to distribution. The structural-buyer paper (2026h) bounds universal ownership and derives the return compression that moderates top accumulation. The counterfactual paper (2026b) compares cohort floors against historical median and mean wealth. The macro paper (2026e) sets the two-circuit structure within which the floor is built. Together they imply a floor that is real but bounded and a top effect that is real but small.

What no paper in the series has done. None models the full distribution — the Gini, the percentile ratios, the top and bottom shares, or inheritance dynamics. The issuance engine tracks a single representative agent per cohort; it has no distribution and so can say nothing about concentration. Measuring the distributional effect requires heterogeneous agents seeded from the real wealth distribution, which is the contribution of this paper.

3. The Model

A microsimulation, not a structural general-equilibrium model. A full heterogeneous-agent general-equilibrium model would be assumption-heavy and its inequality outputs hostage to

unverifiable choices. This paper instead seeds a microsimulation from the actual 2022 SCF household records and applies four explicit channels, each grounded:

1. **The floor.** Each household receives the issuance-engine’s deterministic locked-floor accumulation for its age (per adult), deflated to 2022 dollars. This is the equalising force and the most mechanical channel.
2. **The dividend.** The engine’s standing K3 dividend; mainly income support, with a bounded wealth effect for the portion reinvested.
3. **Return compression.** The structural buyer compresses the realizable return on universally-held assets, moderating top accumulation — bounded by the six-to-twelve-per-cent ownership ceiling of 2026h.
4. **Bequests.** The real distribution determines who holds bequeathable wealth; the transmission is anchored on registry-based causal estimates, not a free-parameter dynamic.

Staging by robustness. Results are reported mechanical-first: the floor’s effect on the bottom (near-assumption-free), then the dividend, then the bounded return-compression band, then the registry-anchored bequest argument. The decomposition asks how much of any Gini change is the floor lifting the bottom versus compression moderating the top.

4. Data and Calibration

The baseline is real, and verified. The seed is the 2022 SCF summary extract (4,595 families, five implicates, full population weights summing to 131.3 million households). Before any CS channel is applied, the weighted records reproduce the published statistics: mean net worth \$1,059,457 (published \$1,059,470), median \$192,700, top-1% threshold \$13.6M, and a wealth Gini of 0.830 — the established US figure. The distribution this paper modifies is therefore the actual one, not a stylised proxy.

Table 1. *The real 2022 SCF baseline the paper starts from (weighted).*

Statistic	Value	Group share of wealth
Wealth Gini	0.830	Bottom 50%: 2.2%
Median net worth	\$192,700	50–90%: 24.4%
Mean net worth	\$1,059,457	90–99%: 38.3%
10th percentile	\$450	Top 1%: 35.1%
Zero or negative net worth	7.9% of households	Under \$10k: 17.6%

Channel calibration sources. The floor is the issuance engine’s base-cohort accumulation by age (the canonical \$209,942 retirement floor, \$190,869 in 2022 dollars), reproduced exactly from the engine. The counterfactual return-wealth gradient is taken from Fagereng, Guiso, Malacrino & Pistaferri (2020), who measure, on Norway’s full population tax records, that the top earns about three percentage points more than the median on financial wealth. The compression is anchored on Gabaix & Koijen (2021), whose inelastic-markets estimate (a multiplier near five, with permanent price impact) makes a structural buyer’s return compression real and lasting, and is bounded in magnitude by 2026h. The bequest channel is anchored on Nekoei & Seim (2023) and Elinder, Erixson & Waldenström (2018).

5. Results

5.1 The Floor: the Robust Result

The bottom of the distribution is transformed. Adding each household's age-appropriate locked floor to the real distribution produces the central, near-mechanical result. The 10th percentile rises from \$450 to about \$64,800. Households with zero or negative net worth fall from 7.9% to about 1.5%. The share of all wealth held by the bottom half nearly triples, from 2.2% to 6.7%. The wealth Gini falls from 0.830 to 0.743.

Table 2. Channel overlay on the real distribution (per-adult floor, 2022\$). The dividend column reinvests in full — an upper bound, since most of it supports consumption.

	Real 2022	+ Floor	+ Floor + Dividend
Wealth Gini	0.830	0.743	0.712
10th percentile	\$450	\$64,800	\$92,300
Median	\$192,700	\$380,500	\$459,300
Bottom-50% share	2.2%	6.7%	8.2%
Top-1% share	35.1%	30.4%	28.6%
Zero/negative NW	7.9%	1.5%	1.0%

Why the top-1% share falls without the top losing anything. The top-1% share declines from 35.1% to 30.4%, but this is almost entirely mechanical: the denominator grew because every household became wealthier, not because the top was diminished. That distinction matters — if the top had collapsed, the result would contradict the bounded-ownership ceiling of 2026h. It does not.

5.2 The Dividend

Mainly income, with a bounded wealth effect. The standing dividend is paid as liquid cash and, for most households, supports consumption rather than accumulation. Reinvested in full it would lower the Gini a further three points (to 0.712) and raise the 10th percentile to about \$92,300, but this is an upper bound; the dividend's primary role is the income floor it provides, which this wealth-focused paper does not fully capture.

5.3 Return Compression: Bounded and Secondary

The most model-dependent channel, stated as such. The structural buyer compresses the realizable return on universally-held assets, slowing top accumulation. The counterfactual divergence is grounded — Fagereng et al. (2020) measure the return-wealth gradient — and the compression is real and permanent (Gabaix & Koijen 2021). But its magnitude on existing private top wealth cannot be observed, because no CS exists; it is bounded by the ownership ceiling of 2026h. Measured over one generation (the horizon over which wealth is held before dispersing at death), the effect is modest:

Table 3. Return compression on the top-1% share over a 30-year generation, by compression magnitude Δ (applied on top of the floor world).

Compression Δ	Top-1% share	Change
----------------------	--------------	--------

0.25 pp	30.4% → 28.9%	-1.5 pp
0.50 pp	30.4% → 27.5%	-2.9 pp
0.75 pp	30.4% → 26.1%	-4.3 pp

The honest reading. Compression trims the top-1% share by one to four points over a generation — real, but secondary to the floor’s near-five-point lift to the bottom half, and the channel whose magnitude rests most on assumption. It restrains the top’s future divergence; it does not pull the top below its current level.

5.4 Bequests: Near-Universal and Permanently Equalising

CS universalises inheritance. Today roughly 76% of households could bequeath something meaningful; under CS, with every citizen holding a locked floor, that approaches 99%. The distribution of inheritances becomes far more equal — the inheritance Gini falls from about 0.83 to 0.74 — because everyone now leaves a floor, not only the already-wealthy.

The registry evidence, and why CS is different. The deeper result comes from the causal literature. Using Swedish and Danish population registers, Nekoei & Seim (2023) and Elinder, Erixson & Waldenström (2018) find that ordinary inheritances reduce relative wealth inequality on receipt — a windfall is large relative to a poorer heir’s prior wealth — but that the equalising effect reverses within about a decade: poorer heirs deplete the windfall while wealthier heirs leave theirs intact and earn higher returns on it. The CS floor bequest is structurally immune to both mechanisms. It is locked, so it cannot be depleted the way liquid cash is; and it earns a uniform return, so there is no rich-heir compounding advantage. CS therefore converts inheritance’s short-lived equalising effect into a permanent one — a conclusion grounded in the very evidence that explains why ordinary inheritance fails to do so.

5.5 Decomposition: Floor versus Top

Overwhelmingly the floor. Of the wealth-Gini reduction from the floor-and-dividend channels, roughly 74% is the floor and at most 26% the dividend (less, since most of the dividend is consumed). The bottom-50% share gains nearly five points from the floor; the top-1% share moves a few points, mostly mechanically, with at most one to four further points from bounded compression. The action is the floor lifting the bottom, not compression at the top. That is the paper’s central decomposition, and it is what keeps the conclusion credible.

6. Robustness

The headline survives the obvious objections. The robust results — elimination of the zero-and-negative tail and the lift to the bottom half — hold across the scenarios a sceptic would raise. Counting one floor per household rather than per adult, allowing the floor to crowd out private saving for households already above it, and shrinking the floor by twenty per cent all leave the qualitative picture intact:

Table 4. *Robustness of the floor result across modelling choices.*

Scenario	Gini	Bottom-50%	Zero/neg
Baseline (no CS)	0.830	2.2%	7.9%
Floor, per adult	0.743	6.7%	1.5%

Floor, one per household	0.764	5.8%	1.7%
Floor with crowding-out	0.748	6.8%	1.5%
Floor 20% smaller	0.758	6.0%	1.8%

Consistency with the structural buyer. A natural objection is that the aggregate floor (\$23 trillion, about 17% of total US net worth) cannot fit within the bounded ownership of 2026h. It does. The floor value is only 19% principal — the structural buyer’s actual purchases, which are what the six-to-twelve-per-cent ceiling bounds — and 81% compound growth accumulated over a career alongside a growing market. The overlay is a mature-CS steady-state comparison; the purchases that build it are bounded exactly as the structural-buyer paper requires.

7. Honest Limits

- The floor result is a steady-state overlay — a comparison of the mature-CS distribution shape with today’s — not a transition path. It answers what the distribution looks like, not how long it takes to get there.
- Return compression is the most model-dependent channel: its inputs are empirical but its magnitude on existing private wealth is a bounded inference, not an observation. The one-to-four-point band reflects that, and the paper does not lean on a point estimate.
- The bequest result rests on registry estimates from Sweden and Denmark; the mechanism (locked, uniform-return) transfers cleanly, but the magnitudes are imported, not re-estimated on US data.
- The dividend’s main effect is on income, not wealth, and is therefore understated here; a companion income-distribution analysis would capture it.
- Behavioural responses — changes in private saving, labour supply, or portfolio choice under a universal floor — are bounded by the crowding-out scenario but not fully modelled.

8. Distributional Claims

Each claim is paired with its limit.

Claim 1 (The floor abolishes the bottom tail). A universal locked floor eliminates zero-and-negative net worth (from 7.9% of households to about 1.5%) and lifts the 10th percentile from \$450 to roughly \$65,000. This is the most robust result and survives every modelling variation tested.

Claim 2 (The Gini falls from below). The wealth Gini falls from 0.83 to about 0.74, with roughly three-quarters of the reduction attributable to the floor lifting the bottom and the top-1% share moving only a few points, mostly mechanically. CS rebuilds the floor; it does not level the top, and 0.74 remains a high-inequality society.

Claim 3 (Top compression is bounded and secondary). Return compression trims the top-1% share by one to four points over a generation, bounded by the ownership ceiling of 2026h. It is real but secondary to the floor and is the paper’s most model-dependent channel.

Claim 4 (CS makes inheritance permanently equalising). By bequeathing a locked, uniform-return floor to every citizen, CS universalises inheritance and — unlike ordinary inheritance, whose equalising effect reverses within a decade (Nekoei & Seim 2023) — makes that equalisation permanent, because the floor can neither be depleted nor compounded faster by the wealthy.

Technical Appendix

A. The Microsimulation

The seed is the 2022 SCF summary-extract public data (4,595 families, five implicates, weight WGT summing to 131.3 million households). All statistics are weight-correct and implicate-correct; the read is validated against the published mean, median, top-1% threshold, and Gini before any channel is applied. The floor by age is the issuance engine's base-cohort balance captured at each age and deflated to 2022 dollars via the engine's own CPI series; the endpoint reproduces the engine's \$209,942 retirement floor exactly. The floor is added per adult (two for married households, one otherwise). Return compression is applied to the top of the floor-adjusted distribution over a 30-year generation. The bequest channel uses the real distribution to determine bequeathable wealth and the registry estimates for transmission. Code and data references accompany the replication package.

Supplementary replication. Three further analyses ground claims made here on the same SCF 2022 microdata, each with parameters swept across cited literature ranges and runnable in the distribution replication package. First, the dividend lane's rent-capitalisation leak is about 1.7% of the dividend centrally (full range 0.5 to 4.6% across Saiz (2010) housing-supply elasticities and Ihlantfeldt low-income demand elasticities, via the standard incidence share), concentrated in the bottom three deciles. Second, the dividend's net demand impulse is small and sign-contested, running from about -0.5% to +0.3% of consumption depending on whether growth-matched issuance displaces bank-created credit (neutral-to-disinflationary) or broadly-held money (mildly inflationary), with a central estimate near +0.2%; the MPC gradient is built from the SCF liquidity distribution and the hand-to-mouth share reproduces the Kaplan-Violante-Weidner one-third benchmark. Third, the floor is about 79% net-new wealth versus 21% crowd-out at the Social-Security-analogue offset (0.33), and remains majority net-new across the full 0.15 to 0.65 pension-displacement literature range; within the bottom five deciles it is 88% net-new, against 73% in the top five. These are first-round, partial-equilibrium bounds rather than general-equilibrium results. The rent and demand-impulse mechanisms follow an approach developed in discussion with the independent researcher wilsonianite, reproduced here on US data.

B. Consistency with the Structural Buyer

The \$190,869 floor (2022\$) decomposes into \$40,727 of principal (the structural buyer's cumulative purchases, 19%) and \$169,216 of compound growth (81%). The six-to-twelve-percent ownership ceiling of 2026h bounds the purchase flow, not the accumulated value; the large value reflects 65 years of compounding alongside a growing market. The aggregate floor of \$23 trillion (about 17% of total net worth) is therefore consistent with bounded ownership in the mature economy.

C. Sensitivity Table

Parameter / scenario	Gini	Bottom-50%	Zero/neg
Floor per adult (central)	0.743	6.7%	1.5%
Floor one per household	0.764	5.8%	1.7%

Floor with crowding-out	0.748	6.8%	1.5%
Floor 20% smaller	0.758	6.0%	1.8%
Floor + dividend (full reinvest)	0.712	8.2%	1.0%

References

- Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109(3), 659–684.
- Alvaredo, F., Garbinti, B., & Piketty, T. (2017). On the share of inheritance in aggregate wealth: Europe and the USA, 1900–2010. *Economica*, 84(334), 239–260.
- Benhabib, J., Bisin, A., & Zhu, S. (2011). The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica*, 79(1), 123–157.
- Boserup, S. H., Kopczuk, W., & Kreiner, C. T. (2016). The role of bequests in shaping wealth inequality: Evidence from Danish wealth records. *American Economic Review*, 106(5), 656–661.
- Elinder, M., Erixson, O., & Waldenström, D. (2018). Inheritance and wealth inequality: Evidence from population registers. *Journal of Public Economics*, 165, 17–30.
- Fagereng, A., Guiso, L., Malacrino, D., & Pistaferri, L. (2020). Heterogeneity and persistence in returns to wealth. *Econometrica*, 88(1), 115–170.
- Feiveson, L., & Sabelhaus, J. (2018). How does intergenerational wealth transmission affect wealth concentration? FEDS Notes, Federal Reserve Board.
- Gabaix, X., & Koijen, R. S. J. (2021). In search of the origins of financial fluctuations: The inelastic markets hypothesis. NBER Working Paper 28967.
- Kuznets, S. (1955). Economic growth and income inequality. *American Economic Review*, 45(1), 1–28.
- Nekoei, A., & Seim, D. (2023). How do inheritances shape wealth inequality? Theory and evidence from Sweden. *Review of Economic Studies*, 90(1), 463–498.
- Piketty, T. (2014). *Capital in the Twenty-First Century*. Harvard University Press.
- Saez, E., & Zucman, G. (2016). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. *Quarterly Journal of Economics*, 131(2), 519–578.
- Federal Reserve Board (2023). *Survey of Consumer Finances, 2022 (Summary Extract Public Data)*.
- Neo-Solon (2026b). *The Citizens Standard: A Historical Counterfactual*. SSRN Working Paper 6735078.
- Neo-Solon (2026e). *The Citizens Standard: A Macroeconomic Model of a Two-Circuit Monetary System*. SSRN Working Paper 6939418.
- Neo-Solon (2026h). *The Citizens Standard: The Structural Buyer*. SSRN Working Paper 6945320.
- Neo-Solon (2026l). *The Citizens Standard: Crisis Behaviour and Failure Modes Under Extreme Conditions*.