THE CITIZENS STANDARD

The Structural Buyer

Asset-Market Dynamics, Price Discovery, and Universal Ownership

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

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

The Citizens Standard: Comparison with Alternative Monetary and Distributional Systems (Neo-Solon, 2026m) · SSRN 6973338

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

Abstract

The Citizens Standard makes the public a permanent, price-insensitive net buyer of total-market equity — on the order of the annual K1+K2 deposit flow. The standing objection, voiced informally by critics and stated formally as an open problem in the macroeconomic model (Neo-Solon, 2026e, §8), is that such a buyer must do one of four things: inflate asset prices without limit, destroy price discovery, leak into consumer-price inflation, or concentrate corporate control in a single state-adjacent vehicle. This paper shows that none of the four holds under the architecture's design, and locates precisely where each effect is real but bounded. The analysis separates the *flow* the buyer adds each period from the *stock* it accumulates over time. On the flow side: the valuation premium converges to a finite fixed point A^*/φ; in steady state the flow is intermediated through primary equity issuance and therefore funds real capital formation rather than perpetual repricing; and the consumption leak into the transactional circuit is bounded by the marginal propensity to consume out of asset wealth (≈ 3%). On the stock side: cumulative ownership plateaus through decumulation rather than rising to one, price discovery is preserved above an active-float threshold, the cost-of-capital effect is allocation-neutral within a total-market index, the index is capture-resistant when mechanical and constitutionally locked, and a mirror-voting rule renders the citizen stake economically real but control-neutral. Each result is conditional on stated quantities; the one that runs against the framework in the present US regime — the net equity-supply response φ — is secured by a buyback constraint this paper identifies as a required complement to the architecture (not yet part of the framework) and carries as its baseline.

Contents

1. Introduction

A monetary architecture that funds citizen wealth by depositing newly issued money into invested savings accounts is, viewed from the capital market, a single agent of unusual character: a buyer of equities that is large, permanent, price-insensitive, and that does not sell on its own account. The natural question — the one a skeptic raises first and the one the macroeconomic model left formally open — is whether an agent of that character breaks the market it operates in.

The objection comes in four forms, and it is worth stating them at their strongest. First, a permanent buyer must bid prices up without limit, so the citizen "wealth" is just an asset bubble the system inflates and then owns. Second, a large price-insensitive holder destroys price discovery, so capital is misallocated because prices stop carrying information. Third, money created to buy equities leaks into the spending economy and shows up as consumer-price inflation, so the dollar of income is debased even as the dollar of wealth rises. Fourth, a single vehicle that accumulates a large equity stake concentrates corporate control, so the architecture quietly nationalizes the means of production.

The reply to all four turns on a single distinction: between the *flow* the buyer adds each period and the *stock* it holds over time. The flow is what could inflate prices, fund or fail to fund real investment, and leak into consumption; it is governed by Propositions 1–3. The stock is what could impair price discovery and concentrate control; it is governed by Propositions 4 and 7. Most of the alarm attached to a "permanent buyer" is in fact alarm about the flow, and the flow results are the most favorable: the premium is bounded, the steady-state flow funds real investment, and the leak is small. The genuinely hard questions live on the stock side, and they are answered there — not by denying that the stock grows, but by showing that decumulation bounds it, that discovery survives above a threshold, and that a voting rule severs ownership from control.

Two features of the architecture do most of the work, and both are worth foregrounding. The first is **two-circuit separation**: the money that funds the floors lives in the asset circuit and does not enter the transactional circuit that prices goods, which is why the asset-side activity does not debase the unit of account. The second is **index breadth**: the Citizens Standard mandates a total-market index, and that single choice minimizes both the cost-of-capital distortion at the index boundary (Proposition 5) and the governance attack surface of index definition (Proposition 6). Breadth does double duty, and the paper returns to it.

This paper closes the structural-buyer problems left open in the macroeconomic model (Neo-Solon, 2026e, §8): the production-side valuation fixed point, the seller-rebalancing channel, and index-definition governance under a permanent buyer. It positions against three literatures it does not otherwise engage: the theory of the universal owner and common ownership (Hawley & Williams, 2000; Azar, Schmalz & Tecu, 2018); the debate over passive investing and price discovery (Grossman & Stiglitz, 1980); and the q-theory of equity issuance and investment (Tobin, 1969; Hayashi, 1982). It takes the two-circuit macro environment of 2026e as given and adds the capital-market layer the series has so far left implicit.

2. The structural buyer, defined

2.1 The agent

The Forced-Deposit Capital Account (FDCA) executes the equity side of the floor system: the K1+K2 deposit flow is directed, each period, into a total-market equity index. Call the net absorption flow A*. At the architecture's calibration A* is approximately $272 billion per year — about 0.39% of US public-equity capitalization (K1 plus the 60 percent of K2 that funds the locked floor under Mode B's 60/40 split; the remaining 40 percent is paid as the standing dividend and is not absorbed into equity) (Neo-Solon, 2026c). In general the flow scales with the adopting economy's aggregates and the single policy lever κ_d (the consumer-dividend share, 0.4 under Mode B): $A^* = (1 - \kappa_d) \cdot g_r \cdot M2 + \kappa_d \cdot K1_{agg}$, or as a share of domestic index capitalization M_{index}, $f = A^*/M_{index} \approx (1 - \kappa_d) \cdot g_r \cdot (M2 / M_{index})$, the approximation holding because the citizenship endowment $K1_{agg}$ is small relative to the growth bulk. A jurisdiction recovers its own figure from its money supply, its index capitalization, and its chosen κ_d — 0.65 percent at $\kappa_d = 0$, 0.39 percent at the Mode B value $\kappa_d = 0.4$, and zero at $\kappa_d = 1$ (Mode D, the all-dividend corner where K1 is also inactive, so the structural buyer absorbs nothing and $A^* = 0$) — without re-deriving the K1 and K2 components. The FDCA is *permanent* (the deposit flow does not stop), *price-insensitive* (it buys to a schedule, not to a valuation signal), and — on its own account — a *non-seller* (it holds; sales occur only through cohort decumulation, §6). It is the cleanest possible structural buyer.

2.2 Scale, in perspective

A buyer of 0.65% of market capitalization per year is not unprecedented in magnitude; it is smaller than the flow the US equity market already absorbs in the opposite direction, where net repurchases have averaged about 1.5% of capitalization per year. The market routinely intermediates a *seller* of larger size. The question this paper answers is therefore not whether the market can absorb the flow, but whether the buyer's permanence, price-insensitivity, and non-selling change the equilibrium in ways the macro model did not resolve.

2.3 Placement in the two circuits

The FDCA operates in the asset circuit M^A. Its purchases are settled in asset-circuit money against equity claims; they do not enter the transactional circuit M^T that prices goods. This is the structural reason the buyer's activity does not, of itself, move consumer prices — the only bridge from M^A to M^T is the consumption leak of §5 (Proposition 3), which is bounded.

2.4 The valuation environment (from 2026e §6)

Let Q_t be the valuation index of the composite equity claim and $Q_{baseline}$ its counterfactual level absent the FDCA. Firms supply equity through the primary market in response to elevated valuations (the supply-side counterpart to the q-theory of investment; Tobin, 1969; Hayashi, 1982), $I_t = \varphi \cdot (Q_t - Q_{baseline})$ with supply elasticity $\varphi > 0$; the tradable float evolves as $\Delta f_t = I_t - A_t$; and valuations adjust to net excess demand for float with speed $\theta > 0$, $Q_{t+1} = Q_t + \theta \cdot (A_t - I_t)$. All parameters are positive. This is the environment in which the flow results are derived.

2.5 The maintained baseline

One calibration fact (developed in §10) is consequential enough to state at the outset. In the present US regime *net* equity issuance is negative — buybacks dominate and rise with valuations — so the supply response φ, taken from current data, is not positive. Proposition 2 is therefore stated under a **buyback-constrained baseline**: a regime in which net issuance responds positively to a sustained valuation premium. This constraint is **now a provision of the Citizens Standard** — the payout-neutrality safeguard of the architecture (Neo-Solon, 2026a, §16.6) and the graduated repurchase charge of the statutory paper (2026d, §10), both built on the net-repurchase base of IRC §4501. The lever is feasible and precedented rather than novel: a 1% repurchase excise was enacted into US law in 2022 and higher rates have since been proposed. Under the buyback-constrained baseline the supply response is the q-theory channel the model assumes; without it, the same flow produces a larger valuation premium and less real investment. The constraint is treated as part of the design this paper proposes, not as an empirical hope.

3. The valuation fixed point (Proposition 1)

Proposition 1. Given a permanent absorption flow A* > 0 and supply elasticity φ > 0, the valuation dynamics admit a unique fixed point **Q* − Q_baseline = A*/φ**, finite for every finite A*, and converge to it iff **0 < θφ < 2** (monotone for θφ < 1, damped-oscillatory for 1 < θφ < 2).

The premium is finite because the supply response is increasing in the premium itself: issuance rises until it equals absorption. It is *decreasing* in φ — a stronger issuance response yields a smaller premium — and diverges only as φ → 0. The structural buyer is dangerous to valuations precisely and only when firms cannot or will not supply equity. *Proof: Appendix A.2.*

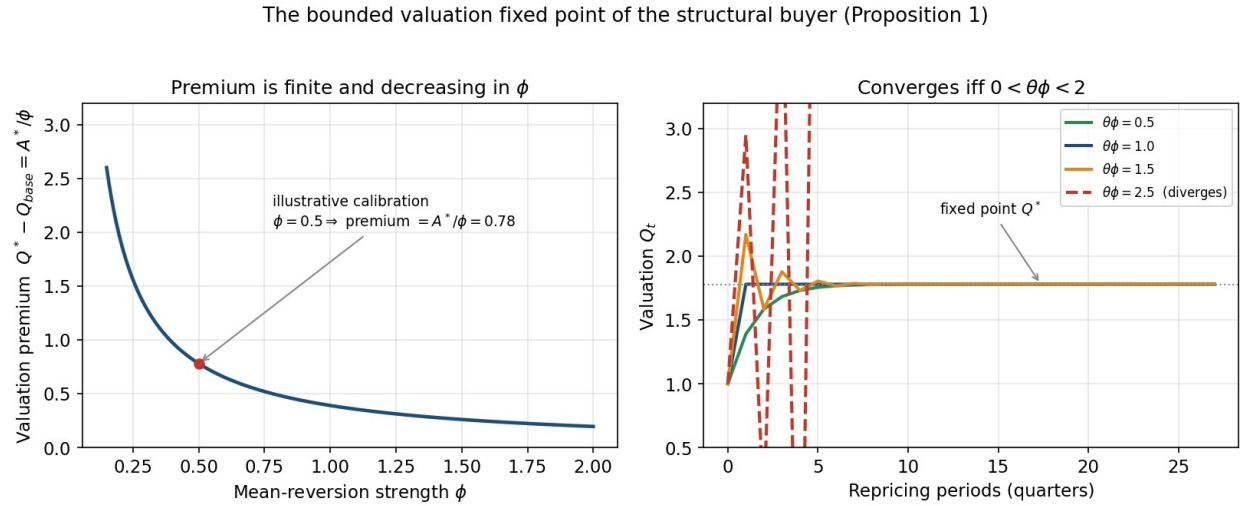

Figure 1. The bounded valuation fixed point (Proposition 1). Left: the valuation premium Q − Q_base = A*/φ is finite for every φ > 0 and decreasing in φ (0.78 at the illustrative φ = 0.5). Right: the quarterly repricing iteration converges to that fixed point if and only if 0 < θφ < 2; at θφ = 2.5 it diverges. The structural buyer raises the level of valuations by a bounded amount without destabilizing price formation.*

4. The flow funds real investment (Proposition 2)

Proposition 2. At Q*, net issuance equals absorption (I* = A*), the float is constant, and firms raise A*·Q* of new equity capital per period through the primary market. The valuation premium A*/φ is a once-for-all level effect realized during the transition; in steady state there is no further repricing. Along the transition, per-period repricing decays geometrically at rate |1 − θφ| while the share of the flow met by primary issuance rises monotonically from 0 to 1, with total repricing equal to A*/φ.

This is the formal answer to "it just inflates assets." In steady state the perpetual flow is intermediated through primary issuance and funds real capital formation; the premium is a one-time repricing, not an ongoing inflation. The result rests on reading φ as a genuine equity-issuance/investment elasticity — the maintained buyback-constrained baseline of §2.5 — and that interpretation is the paper's central empirical commitment and falsification point (§10). *Proof: Appendix A.3.*

5. Seller rebalancing and the consumption leak (Proposition 3)

Proposition 3. For any FDCA secondary-market purchase Δ, the consumption leak into the transactional circuit is bounded above by κ_W·Δ, where κ_W ∈ [0.025, 0.05] is the marginal propensity to consume out of asset wealth; hence at least (1 − κ_W) of secondary proceeds re-enters the asset circuit.

The seller of equity rebalances toward a target money share; the only escape to the goods-pricing circuit is consumption out of the realized wealth change, governed by κ_W. With κ_W near 0.03 (§10.3) the leak is at most a few percent — and, by Proposition 2, in steady state the FDCA's counterpart is an *issuing firm investing the proceeds*, not a household consuming them, so the leak bears chiefly on the transition. *Proof: Appendix A.4.*

6. Price discovery and the stock problem (Proposition 4)

6.1 The stock problem

The FDCA never sells on its own account, so taken in isolation its ownership share ψ_t rises without bound — a permanent non-selling buyer eventually owns everything. If that were the whole story, price discovery would fail for want of a tradable float.

6.2 Decumulation bounds ψ

The premise is false at the system level. The Stable Floors fund retirement; as each cohort decumulates, its floor is liquidated and equity returns to the active float. The FDCA is a permanent gross buyer with a demographically-lagged gross seller attached. Modeling the aggregate floor as a funded stock gives a steady-state ownership share **ψ* ≈ c · dur** (deposit-rate-in-equity-terms times average holding duration; the stock-flow identity is derived in Appendix A.6). In the base calibration (c ≈ 0.004, *dur* ≈ 30 years) **the zero-growth value is c·dur ≈ 0.12 (≈ 0.16 at the cohort-realistic dur ≈ 40 years); under the framework's 2 percent real growth the realized share is ψ* = c·annuity(g, dur) ≈ 0.10 (the constant-hazard and fixed-duration decumulation models bracket 0.09–0.11)**, leaving ~88 percent

active float; the no-decumulation counterfactual diverges past one. The universal-owner concern (Hawley & Williams, 2000), and the related common-ownership concern (Azar, Schmalz & Tecu, 2018), thus bear on a *bounded* 10–20 percent holding — the realized Mode B share (≈10%) up to its c/g ceiling (≈20%) — not on literal universal ownership. These figures are the US calibration; the load-bearing bound is not the ten-percent point estimate but the c/g ceiling itself, and it moves with the economy the framework runs in. Because the ceiling is the flow-to-growth ratio, a shallower equity market or a faster-growing economy raises it: India, the fastest grower in our calibration set and with a low transaction-active money share, yields a realized share near forty-seven percent against a c/g ceiling near fifty-seven percent, rather than the US ten and twenty. The bound remains binding and well inside the feasibility limit in every case, but the specific ownership share is a property of the implementing economy, not a universal constant.

6.3 The threshold

Informativeness depends on the active price-setting margin, not on the size of a non-trading holding — the Grossman–Stiglitz logic that prices stay informative as long as informed trading is rewarded (Grossman & Stiglitz, 1980): a buyer that holds and mirror-votes removes depth rather than injecting noise. Discovery is preserved while active float $(1 - \psi_t)$ exceeds a threshold f_{min} and degrades continuously, not catastrophically, toward it. Estimates of US passive ownership span ~18% of the market on narrow fund-holdings measures (ICI) to ~33–41% on demand-elasticity and reconstitution-volume measures (Haddad, Huebner & Loualiche, 2025; Chinco & Sammon, 2024) — and price discovery continues to function across that range, which suggests f_{min} sits below the conservative (zero-growth) active float of 85%. That the estimates themselves vary more than twofold is precisely why a backstop is carried.

Supplementary replication. A cohort accumulation-and-decumulation simulation tests this plateau directly and largely confirms it. Starting from the verified absorption flow (≈ 0.39 percent of market capitalization per year at Mode B) and releasing each cohort's holdings back to the float after a working-life holding period, the ownership share converges in every parameterization and never trends toward one — the decumulation bound of Proposition 4 holds structurally, and the no-decumulation counterfactual diverges as expected. The simulated plateau reproduces the closed-form $\psi^* = c \cdot \text{annuity}(g, \text{dur})$ to within a fifth of a percentage point, and the central case (0.39 percent flow, a ~40-year holding duration, 2 percent growth) settles at 10.9 percent — inside the stated 0.09–0.11 band. The plateau *level* is duration- and flow-sensitive: across plausible holding durations (25–50 years) and flow rates the plateau ranges roughly 6 to 21 percent, with the ~10 percent central case settling at the verified flow and a moderate duration and the high-teens reached at long durations or heavy floor-weighting. In every case the active float remains a clear majority (~79–94 percent), well above the empirical range where price-discovery concerns begin to bind, so the float-threshold result holds throughout. Runnable as the structural-buyer-endgame module in the distribution replication package.

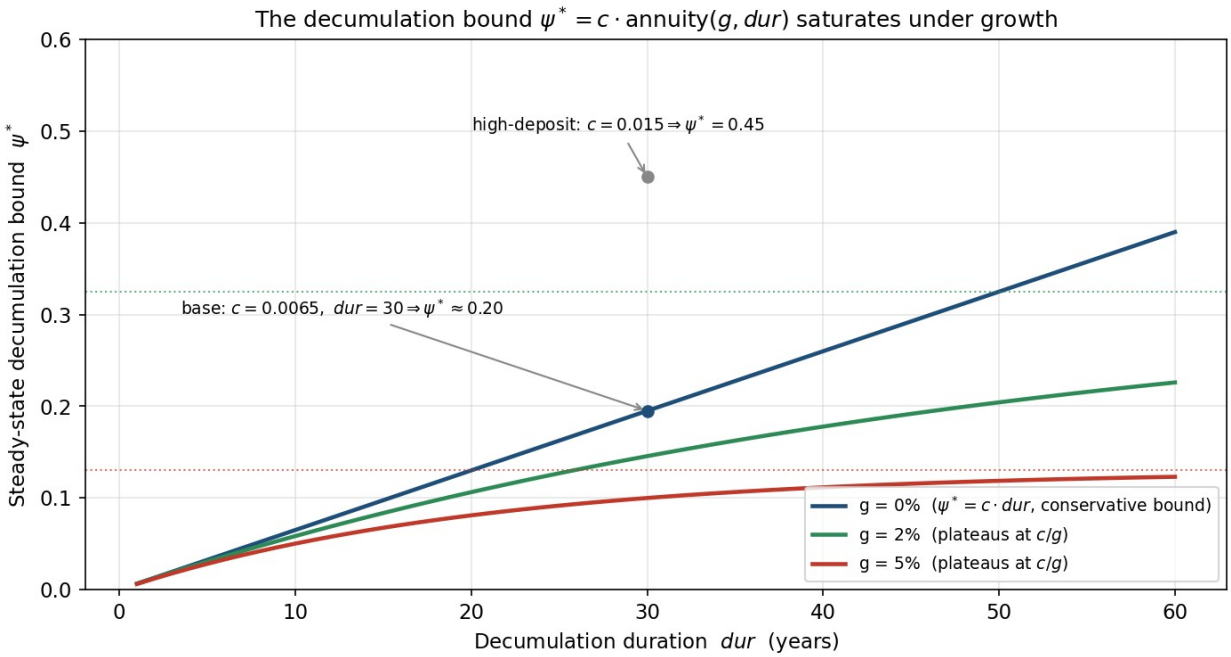

Figure 2. The decumulation bound ψ = c · annuity(g, dur) saturates under growth. At zero growth ψ* = c · dur, the conservative bounds (maximum-absorption κ_d = 0: c = 0.0065, dur = 30 ⇒ 0.20; high-deposit: c = 0.015 ⇒ 0.45). Under positive growth ψ* rises more slowly and plateaus at c/g, so the steady-state public ownership share stays bounded rather than climbing toward one.*

6.4 Proposition 4

Under decumulation recycling, ψ converges to a finite ψ* ≈ c·*dur* rather than to one; price discovery is preserved while 1 − ψ* exceeds f_min and degrades continuously as it approaches f_min.

6.5 Design decision — the holdings ceiling

Because ψ* rises with the deposit rate and duration, the architecture carries an explicit **constitutional holdings ceiling ψ_max < 1 − f_min**. On approaching it, the marginal deposit is redirected from the saturated domestic index into a broader asset universe (international equity, then a wider claim set) and/or float recycling is accelerated. The ceiling governs liquidity and discovery only — by Proposition 7 control is neutralized at every ψ — and it is a backstop, not the primary mechanism: in the base case it never binds. Redirection on saturation is where the external paper (2026g) enters.

A third redirect on saturation is internal to the citizen circuit and, unlike broadening the buy-side, keeps the diverted purchasing power in citizens' hands. The marginal deposit that would breach the holdings ceiling is moved from the locked index-purchase channel (K2) to the liquid dividend channel (KI). This redirect is price-neutral by construction: Mode Ω (Architecture, Section 8) re-solves the K2 and KI mix so that derived consumer-price inflation stays at zero as the mix shifts, because lowering K2 opens a deflationary gap that KI fills exactly. The holdings ceiling therefore becomes a tunable, price-stable dial rather than a hard backstop: any target ownership share at or below the natural level is attainable at zero price drift, with the value the citizen would have held as equity reappearing as spendable dividend. It is a change in the form of citizen benefit, ownership for income, not a reduction in it.

The cost of maintaining a target share is smaller than it first appears, and it connects directly to Proposition 2. To hold a steady ownership share X, the citizen circuit need only absorb X of net new equity issuance N each period: the citizen stake and the market both appreciate at the same return, so appreciation cancels in the share and only new stock the citizen did not buy dilutes it. The required absorption is thus $A = X \cdot N$, which at the full-ownership limit $X = 1$ recovers Proposition 2's fixed point $I^* = A^*$ (absorption equals net issuance, constant float). Because net issuance is a small flow relative to market capitalization in most markets (for India, net equity issuance of roughly 0.66 percent of market capitalization per year against real growth of 6.5 percent), holding even a large ownership share imposes only light absorption, and the ceiling tunes it downward from there. In a net-buyback regime such as the recent United States, where N is negative (Section 10.2), the same relation implies the citizen share rises without any purchase and is held steady by taking distributions rather than buying, a materially different regime treated in the external paper (2026g).

7. Cost of capital and capital allocation (Proposition 5)

Proposition 5. A float-weighted structural buyer is allocation-neutral within the index — it lifts each member's valuation in proportion to the weight the market already assigns, leaving relative valuations unchanged — and creates an inclusion-boundary premium $\sigma \leq A^*/\varphi$ across the index edge that is decreasing in index breadth.

Neutrality is *relative*, not corrective: the buyer preserves the prevailing market-cap allocation proportionally, neither disrupting nor repairing it, so the honest claim is only that no *new* relative distortion is introduced inside the index. The boundary premium σ is minimized by construction, because the Citizens Standard mandates a total-market index: as coverage approaches the whole listed universe, the excluded set shrinks to the micro-cap tail and $\sigma \to$ small. The index-addition literature both bounds σ and confirms its direction: the S&P 500 inclusion premium fell from about 7.4% in the 1990s to under 1% over the past decade, even as indexed assets grew (Greenwood & Sammon, 2025; cf. Patel & Welch, 2017) — small and shrinking exactly as breadth rises, which is both the empirical content of Proposition 5 and its direct falsification test.

8. Index governance and capture resistance (Proposition 6)

Proposition 6. A mechanically float-weighted, objectively-thresholded, committee-free index, amendable only by Tier-2 constitutional supermajority, admits no profitable manipulation of inclusion or weighting; the residual threshold-gaming incentive is bounded by σ and minimized under total-market coverage.

The attack surface in index design is discretion: where a committee chooses constituents, inclusion becomes worth lobbying for and selling. Removing discretion removes the lever — a firm is included iff it meets public, mechanical criteria satisfiable only by real economic action — and the constitutional lock raises the cost of changing the rule above any inclusion rent. The residual incentive to game a mechanical threshold is bounded by the inclusion premium σ, which §7 shows is small under total-market coverage: when almost everything is included, there is little rent in inclusion. This is an incentive-compatibility argument, conditional on the credibility of the constitutional lock that the transition paper's lock-credibility model (2026c) is built to establish.

9. Corporate control: ownership without control (Proposition 7)

Proposition 7. If the FDCA votes its shares in the same proportions as the aggregate of all non-FDCA shareholders (mirror-voting), then for every corporate vote the outcome is identical to that determined by the non-FDCA shareholders alone: the FDCA bloc is never pivotal. The citizen stake conveys full cash-flow rights and zero marginal control.

The proof is a one-line identity — mirror-voting adds a positive scalar multiple of the residual vote to itself, which cannot change which side holds the majority (Appendix A.5). The rule dominates mandated abstention, which fails outstanding-share thresholds by counting the FDCA's stake effectively as "no" and thereby handing it a *de facto* veto as ψ grows. Mirror-voting is what licenses the paper's central claim: universal economic ownership need not imply concentrated corporate control.

10. Calibration and falsification

The propositions are conditional on five quantities. Four are favorable or neutral; the fifth runs against the framework in the present regime and is secured by design.

10.1 A* (absorption flow)

≈ $272B/yr ≈ 0.39% of market cap (Neo-Solon, 2026c) — smaller than current net buybacks (~1.5%). Neutral.

10.2 φ (net equity-supply response) — the binding constraint

Governs Propositions 1–2. In the current US regime net issuance is negative — about −$580B in 2016, and negative for all but one of the last fifteen years for nonfinancial firms (Palladino, 2019) — and buybacks are procyclical, so φ from current data is not positive — the opposite sign to Proposition 2's requirement. The *gross* issuance margin behaves as assumed (firms raise equity when valuations are high; Baker & Wurgler, 2002; Graham & Harvey, 2001), so the channel exists; it is swamped by repurchases. Resolution: a **buyback constraint** (graduated repurchase tax or cap). This lever is **a provision of the Citizens Standard**, set out as the graduated repurchase charge of the statutory paper (Neo-Solon, 2026d, §10) and the payout-neutrality safeguard of the architecture (2026a, §16.6). It builds directly on existing law — the 2022 Inflation Reduction Act's 1% excise on net repurchases (IRC §4501, whose base already nets issuance against repurchases), with a 4% rate proposed and graduated higher here so it binds. The magnitudes are not marginal: S&P 500 buybacks set a record $942.5B in 2024 (~1.4% of total-market cap) and net nonfinancial equity issuance has been negative every year since 1997 (Palladino, 2019), several times the structural buyer's ~0.39%. The empirical record sharpens the design rather than merely supporting it: the 2022 excise cut repurchases by roughly a quarter without an offsetting rise in dividends, raising cash but not investment (Autore, Barnes, Clarke & Schrowang, 2025), while the issuance response to high valuations is robust in aggregate yet concentrated among firms with an existing need for capital (DeAngelo, DeAngelo & Stulz, 2010). A weak constraint therefore stops buybacks without redirecting to investment; the permanent bid must do the issuance-pulling, and it pulls hardest where fundable opportunities exist. Constraint and premium are jointly necessary — neither alone produces Proposition 2's result. The paper carries the buyback-constrained regime as baseline.

Falsification: impose a sustained premium and measure net issuance; if it fails to turn positive under a binding constraint, Proposition 2 fails and the premium is large. *Adverse — secured by lever, not assumed.*

10.3 κ_W (MPC out of asset wealth)

Governs Proposition 3. Modern estimates ≈ 3.2¢/$ (Chodorow-Reich, Nenov & Simsek, 2021); the aggregate propensity to consume out of wealth has declined since the Great Recession, pushing it lower still (Beach, Gamber & Moran, 2025); equity-specific ≈ 3¢ (Poterba, 2000). Range [0.025, 0.05], central ≈ 0.03 — and plausibly lower for locked, illiquid floors. *Favorable; leak ≈ 3%.*

10.4 θ (adjustment speed)

Governs Proposition 1 stability. Quarterly repricing; only $0 < \theta\varphi < 2$ is required, easily met for modest φ. Estimate jointly with φ from the transition path. *Neutral.*

10.5 μ* (target money share)

Governs Proposition 3 dynamics, not the headline bound. Household liquid-asset shares anchor it. Neutral.

10.6 V_T stability (M^T tracks Y)

Governs the price-stability result. Aggregate velocity is unstable, but the two-circuit claim is narrower and testable: transactional velocity V_T should be *more* stable than aggregate velocity because asset-driven money demand is sequestered in M^A. *Open — testable; add a velocity-feedback term to the issuance rule if it fails.*

11. Limitations and scope

This paper isolates the capital-market layer and holds several things fixed that belong elsewhere. **Behavioral responses** — labor supply under guaranteed wealth, consumption-smoothing against floor balances — are not modeled; they belong in a micro-foundations treatment. **Crisis-state and drawdown dynamics** — what a market crash does to floors near maturity, the sequence-of-returns problem — are deferred to the banking paper's stress machinery (2026f) and the planned decumulation paper; the price-discovery and control results here are steady-state and do not characterize tail states. **International equity and home bias** enter only as the destination of redirected deposits at the holdings ceiling (§6.5) and are developed in the external paper (2026g). The **high-ψ corner**, where ψ* approaches 1 − f_min under aggressive deposit or adverse demographic calibration, is handled by the ceiling but not characterized in detail. Two dependencies are load-bearing and stated as such: Proposition 2 requires φ > 0, secured by the buyback constraint (§10.2), and Proposition 6 requires the constitutional lock to be credible (2026c). The price-discovery threshold f_min is taken from an unsettled literature; the paper's claims degrade gracefully in f_min rather than depending on a sharp value.

12. Conclusion

The structural buyer is the point at which the Citizens Standard touches real capital markets most forcefully, and the point at which its critics press hardest. Separating the flow from the

stock dissolves most of the force of the objection. On the flow side, the valuation premium is bounded at A*/φ, the steady-state flow funds real investment through primary issuance rather than inflating prices, and the consumption leak is a few percent — provided the buyback-constrained baseline holds, which is a design choice this paper adds to the architecture rather than one it already contains. On the stock side, decumulation bounds cumulative ownership to a plateau near 10 percent of the market under the framework's Mode B calibration — rising toward a quarter to a half only under maximum-absorption or high-deposit assumptions — rather than to universal ownership, price discovery survives above a float threshold backstopped by a constitutional holdings ceiling, the cost-of-capital effect introduces no new relative distortion, the index is capture-resistant when mechanical and locked, and mirror-voting severs ownership from control completely. The architecture's two recurring features — two-circuit separation and total-market breadth — are doing the load-bearing work throughout. What remains genuinely open is empirical, not conceptual: the sign and size of the net-supply response under a binding buyback constraint, and the stability of transactional velocity. Both are falsification points, stated as such. The paper closes the structural-buyer problems the macroeconomic model left open and hands the next paper — decumulation and sequence-of-returns risk — the float-recycling mechanism on which its own ψ-plateau result depends.

Appendix A — Formal statements and proofs

Proofs of the four analytically tractable propositions (1, 2, 3, 7) and the ownership steady state of Proposition 4. Each was confirmed numerically before being written: the fixed point and steady-state identity to machine precision, the stability band convergent at θφ ∈ {0.5,1.0,1.5} and divergent at 2.5, and the mirror-voting identity across the full (p, ψ) grid.

Setup. $Q_{t+1} = Q_t + \theta(A_t - I_t)$, with $I_t = \varphi(Q_t - Q_{baseline})$, $A_t = A^* > 0$, $\theta, \varphi > 0$.

A.2 (Proposition 1). Substituting, $Q_{t+1} = (1 - \theta\varphi)Q_t + \theta(A^* + \varphi Q_{baseline})$: an affine map $Q_{t+1} = aQ_t + b$ with $a = 1 - \theta\varphi$, $b = \theta(A^* + \varphi Q_{baseline})$. The unique fixed point solves $Q^*(1-a) = b$; since $1 - a = \theta\varphi \neq 0$, $Q^* = b/(\theta\varphi) = Q_{baseline} + A^*/\varphi$, finite because $\varphi > 0$. The general solution $Q_t = Q^* + a^t(Q_0 - Q^*)$ converges iff $|a| < 1$, i.e. $0 < \theta\varphi < 2$; $a \in (0,1)$ gives monotone convergence, $a \in (-1,0)$ damped oscillation. ∎ *Remark:* $\partial(Q^* - Q_{baseline})/\partial\varphi = -A^*/\varphi^2 < 0$; the premium diverges only as $\varphi \to 0$.

A.3 (Proposition 2). At Q^*, $I^* = \varphi(Q^* - Q_{baseline}) = \varphi(A^*/\varphi) = A^*$, so $\Delta f^* = I^* - A^* = 0$ and primary capital raised is $I^* \cdot Q^* = A^* \cdot Q^*$ per period; Q constant ⇒ repricing $R_t = 0$ in steady state. On transition from $Q_0 = Q_{baseline}$, $R_t = Q_{t+1} - Q_t = a^t(a-1)(Q_0 - Q^*) = a \cdot R_{t-1}$, geometric decay at $|a| = |1 - \theta\varphi|$. The issuance share $I_t/A^* = (Q_t - Q_{baseline})/(Q^* - Q_{baseline})$ rises monotonically $0 \to 1$, and $\Sigma R_t = Q^* - Q_0 = A^*/\varphi$. ∎ *Remark:* rests on $\varphi > 0$ as a real issuance/investment elasticity (q-theory); if proceeds were returned via buybacks, $\varphi \approx 0$ and the premium is larger — the §10.2 falsification point.

A.4 (Proposition 3). A seller with $W = m + e$ at optimum money share $\mu^* = \mu(r, \eta)$ is swapped Δ of equity for Δ of money; in a pure swap W is unchanged, and at a premium the realized gain is $\Delta W \leq \Delta$. Rebalancing to μ^* re-deploys money to equity (back to M^A); the only escape to M^T is consumption $\kappa_W \cdot \Delta W \leq \kappa_W \cdot \Delta$. With $\kappa_W \in [0.025, 0.05]$, leak ≤ 5%, re-deployment ≥ 95%. ∎ *Remark:* by A.3 steady-state purchases are primary, so the leak is chiefly transitional; $\kappa_W \cdot \Delta$ is a conservative bound.

A.5 (Proposition 7). FDCA holds share ψ; residual shareholders hold (1 − ψ) and cast fraction p YES. Under mirror-voting the FDCA casts ψp YES, ψ(1−p) NO. Total YES = (1−ψ)p + ψp = p, independent of ψ; for any threshold τ the proposal passes iff p ≥ τ — the residual outcome. Hence the FDCA is never pivotal. ∎ *Corollary:* abstention yields YES-among-cast = p but, on outstanding-share thresholds, counts ψ as effective NO, giving a de facto block that grows in ψ; mirror-voting preserves p under both threshold types and so weakly dominates, strictly on outstanding-share thresholds.

A.6 (Proposition 4, ownership steady state). The FDCA acquires equity at a steady per-period rate and releases it through cohort decumulation. Let c be the steady-state net acquisition per period expressed as a share of market capitalization (c ≈ A*/mktcap ≈ 0.39%/yr), and let each unit of acquired equity remain in the floor system an average holding duration *dur* before decumulation liquidates it back to the active float. The citizen equity share is reported on this same basis: ψ_t = (cumulative net FDCA acquisitions) ÷ (total market capitalization), the public's running ownership stake in the index. Market capitalization enters the framework only as this reporting denominator — the issuance rule itself is anchored to M2 and the transactional circuit M^T (Neo-Solon, 2026e), never to the market's valuation, so a rising or falling market changes the reported share without altering a single dollar of scheduled issuance. In demographic steady state the FDCA-held stock is constant when acquisition equals release. Modeling release as a constant hazard 1/*dur* (mean holding time *dur*), steady state requires inflow = outflow: c = ψ*/*dur*, hence **ψ* = c·dur**. Equivalently, in cohort form, current holdings are the unliquidated acquisitions of the last *dur* periods, Σ_{k=0}^{*dur*−1} c = c·*dur*. This is the stock-flow (Little's Law) identity: a steady-state stock equals arrival rate times average residence time. With c ≈ 0.004 and *dur* ≈ 30 (≈40 on the cohort-realistic estimate), the zero-growth value is c·dur ≈ 0.12–0.16 and the realized share under 2 percent growth is ψ* ≈ 0.10; the no-decumulation limit *dur* → ∞ gives ψ* → ∞, so decumulation is exactly what bounds ψ. ∎ *Remark:* the result is a steady-state identity, not a behavioral assumption — only the values of c and *dur* are empirical (Appendix B), and ψ* scales linearly in each. It is leading-order: cohort-duration heterogeneity and non-stationary market-cap growth perturb it at second order (in the value-share interpretation, positive growth lowers it, since older cohorts were acquired into a smaller market), so c·*dur* is a conservative estimate of a plateau that is in any case finite and well below one.

Appendix B — Calibration tables

B.1 Parameters.

Symbol	Meaning	Working value/range	Source	Governs
A	absorption flow	0.39–0.65% of mkt cap/yr	architecture calibration	Props 1–2
φ	net equity-supply response	> 0 under the repurchase constraint (2026d §10)	§10.2 (BIS; Palladino 2019)	Props 1–2
θ	valuation-adjustment speed	s.t. 0 < θφ < 2	transition path	Prop 1
κ_W	MPC out of asset	0.025–0.05 (≈0.03)	Chodorow-Reich et	Prop 3

Symbol	Meaning	Working value/range	Source	Governs
	wealth		al. 2021; Beach et al. 2025	
μ	target money share	household liquid share	balance-sheet data	Prop 3 dyn.
ψ	steady-state ownership	≈ 0.10 realized (c·dur ≈ 0.12–0.16 at zero growth)	§6.2 funded-stock model	Prop 4
f_min	discovery threshold	< 0.6 (tentative)	Chinco-Sammon 2024; Haddad et al. 2025	Prop 4
σ	inclusion-boundary premium	≤ A/φ; <1% recent decade	Greenwood-Sammon 2025; Patel-Welch 2017	Prop 5

B.2 Premium comparative statics (illustrative). Premium = A*/φ, decreasing in φ; a one-time level shift realized over the transition, with per-period repricing decaying at |1 − θφ|. Over the supply-elasticity bracket φ ∈ [0.5, 2.0] (2026e §6.5): in the accumulation phase, where absorption runs at the gross deposit rate A* ≈ 0.39% of capitalization, the transitional peak premium is ≈0.78% (φ=0.5), ≈0.39% (φ=1.0), ≈0.20% (φ=2.0) of valuation; at demographic maturity, where drawdown nets absorption down to A* ≈ 0.16% (2026e §6.5), the standing premium is ≈0.32% / 0.16% / 0.08% correspondingly. Either way the premium is a few tenths of one percent of valuation and eases as net absorption falls with demographic maturity — consistent with the macro paper's worked bound.

B.3 ψ plateau (illustrative). ψ* ≈ c·dur: Mode B (c≈0.004, dur=30) → 0.12 (≈0.10 under 2% growth); maximum-absorption κ_d=0 (c=0.0065) → 0.20; high-deposit (c=0.015) → 0.45; no-decumulation counterfactual → diverges past 1. Active float 1 − ψ* = 0.90 (Mode B), 0.80 (max-absorption).

Appendix C — Price-stable ownership targeting

This appendix derives the result summarized in Section 6.5: under Mode Ω the citizen ownership share can be set to any target at or below its natural level while the consumer price level is held exactly flat, and the purchase required to hold a target share is a share of net new equity issuance rather than of total market growth. The construction generalizes Proposition 2: at the full-ownership limit it recovers the fixed point in which absorption equals net issuance.

C.1 Setup

Write the framework's derived-inflation identity for one economy in one period. Let μ be the transaction-active money share, m2 broad money, gdp output, mkt equity market capitalization, rg the real growth rate, and NC the net citizen-accrual constant. The four issuance-channel parameters are k1, k2, k3, ki. Define:

```
MT = μ·m2,    k1Agg = k1·gdp·NC,    S = max(0, k2·m2·rg − k1Agg)
  cpi = ( k3·S + ki·m2 + 0.20·(k1Agg + (1−k3)·S) ) / MT − rg
```

S is the growth-funded budget after the K1 endowment. The split sends k3·S to the citizen dividend and (1−k3)·S into the locked floor, which purchases the broad index. Two derived quantities matter: price stability (cpi = 0), and the ownership accrual O = (1−k3)·S, the floor's annual index purchase.

C.2 The steady-share condition

The citizen stake and the total market both appreciate at the same equity return r_p, because the locked floor holds the broad index. Over one period:

$$\text{stake}' = \text{stake}\cdot(1 + r_p) + O, \qquad \text{mkt}' = \text{mkt}\cdot(1 + r_p) + N$$

where N is net new equity issuance (gross issuance minus buybacks). Imposing a constant share stake/mkt = X in both periods, the appreciation terms cancel and the condition reduces to a single equation:

$$O^* = X\cdot N$$

To hold a share X, the citizen circuit must absorb X of each period's net new issuance, no more. Appreciation accrues to citizen and market alike and never enters. This is exact, and N is a directly measured flow, so no market-growth proxy is required.

C.3 Proposition C.1 (existence and uniqueness)

Proposition C.1. Fix an economy's calibration and an ownership target X at or below the natural level. Holding the split at Mode Ω's price-stabilizing value k3*, there exists a unique pair (k2, ki) with ki ≥ 0 such that the configuration (k2, k3*, ki) satisfies both cpi = 0 and O = X·N.

Proof. The system is triangular. The ownership condition fixes the budget:

$$(1-k3^*)\cdot S = X\cdot N \quad \Rightarrow \quad S^* = X\cdot N \,/\, (1-k3^*)$$

Since S = k2·m2·rg − k1Agg is monotone in k2, S* fixes k2 uniquely:

$$k2^* = \left(S^* + k1Agg \right) / \left(m2\cdot rg \right)$$

With S* and k3* fixed, the price-stability equation is linear in the single remaining unknown ki and solves in closed form:

$$ki^* = \left[rg\cdot MT - k3^*\cdot S^* - 0.20\cdot(k1Agg + (1-k3^*)\cdot S^*) \right] / m2$$

Ownership fixes S hence k2; price stability then fixes ki. The solution is unique and admissible whenever ki* ≥ 0. Lowering the target lowers S*, which lowers total issuance and pushes derived inflation downward, so the residual gap ki must fill is deflationary and ki* rises above zero. Hence every target at or below the natural level (the level produced when ki = 0) is admissible. □

Interpretation of the boundary. KI adds money but cannot withdraw it; the construction only ever asks it to add, offsetting the deflationary gap created by reducing index purchase. The one direction the lever cannot go is above the natural ownership level, which would require net contraction (the Surge Brake, not KI). As a ceiling the mechanism is therefore always available.

C.4 Relation to Proposition 2

At X = 1 the steady-share condition O* = X·N becomes O* = N: the citizen circuit absorbs all net new issuance, which is exactly the fixed point of Proposition 2 (absorption equals net issuance, I* = A*, constant float). Proposition C.1 is thus the partial-share generalization of Proposition 2: the full-ownership case is one point on a continuum of price-stable ownership targets, each with its own (k2, ki).

C.5 Worked calibration: India

Primary-source inputs (RBI M3 and time-deposit share for μ; MoSPI GDP; NSE market capitalization), with India's measured net equity issuance:

```
μ = 0.28,  m2 = 281.4T,  gdp = 330.7T,  mkt = 438.9T,  rg = 6.5%
MT = 78.79T,  k1Agg = 0.097T,  Ω split k3* = 10.05%,  N ≈ 2.9T/yr
```

The issuance figure is built from 2024 primary-market data: IPOs raised about 1.34T, QIPs a record 1.37T, with rights and follow-on offerings bringing gross issuance to roughly 3.0T, against buybacks of only about 0.13T (which collapsed after the October 2024 shift of buyback tax to shareholders). India is a strong net issuer at N ≈ 2.9T, about 0.66 percent of market capitalization per year. India runs +5.53 percent derived inflation under a fixed 60/40 split; Mode Ω's split solve (k3* = 10.05 percent, matching the closed form κ_d* ≈ (μ − 0.20)/0.80) brings that to zero. Applying the ownership-targeting construction on top, holding cpi = 0 throughout:

At X = 52.5 percent: annual purchase O* = 1.52T, k2* = 0.098, ki* = 1.644 percent. At X = 30 percent: 0.87T, k2* = 0.058, ki* = 1.717 percent. At X = 15 percent: 0.44T, k2* = 0.032, ki* = 1.765 percent. At X = 5 percent: 0.15T, k2* = 0.014, ki* = 1.797 percent. In every row derived inflation is 0.000 percent and ki* is positive. Holding even a 52.5 percent share requires purchasing only about 1.5T a year, because the float grows by new issuance of only 0.66 percent annually, not by the 6.5 percent real growth rate; a growth-based estimate overstates the required purchase by roughly an order of magnitude.

C.6 The net-buyback regime

Where net issuance is negative, as in the recent United States (Section 10.2), an index holder's share rises automatically as the float shrinks, so O* = X·N < 0: holding a steady share requires net selling, taking distributions rather than buying. The construction still applies, but the citizen circuit takes cash out rather than adding to the stake, a materially different regime. This is the setting of the external-interoperability treatment (Neo-Solon 2026g).

C.7 Scope

The result is a theorem about the framework's derived-inflation identity; it inherits whatever gap exists between that identity and realized inflation and does not independently validate it. The inputs are primary-source empirical; the mechanism and price path are model quantities. The desirability of any particular ceiling is a policy choice, not a result of the theorem, which establishes only that any target at or below the natural level is attainable at zero price drift. Maintaining a target share requires re-solving against each period's measured net issuance; a multi-period path with a time-varying issuance series is the natural empirical extension. A final scope condition concerns the issuance floor that bounds Mode Ω generally: because a lower ownership ceiling reduces the growth-funded budget S, a sufficiently low target can in principle press S toward the non-negativity floor (S ≥ 0) that the mode already requires. Whether this binds is economy-specific and depends on the size of the K1 endowment relative to net issuance. For the India calibration it does not bind anywhere in the range considered: since the ownership budget (1−k3*)·S = X·N is non-negative, k2 is bounded below by k1Agg/(m2·rg) ≈ 0.005 and stays strictly positive even at a five percent ceiling, so the construction remains admissible throughout. Economies with a large K1 endowment relative to their net equity issuance should nonetheless verify the floor is clear before applying a very low ceiling.

References

Neo-Solon (2026a–g). *The Citizens Standard*, Papers 1–7 (Architecture; Counterfactual/Empirical; Transition; Statutory; Macroeconomic Model; Banking; External Interoperability). SSRN 6702518; 6735078; 6810741; 6873798; 6939418; 6939498; 6939600. Replication: github.com/Neo-Solon/Citizens-Standard.

Autore, D. M., Barnes, S., Clarke, N., & Schrowang, A. (2025). Corporate share repurchases and the 2023 excise tax. Journal of Corporate Finance.

Azar, J., Schmalz, M. C., & Tecu, I. (2018). Anticompetitive effects of common ownership. Journal of Finance, 73(4), 1513–1565.

Baker, M., & Wurgler, J. (2002). Market timing and capital structure. Journal of Finance, 57(1), 1–32.

Bank for International Settlements (2020). BIS Quarterly Review, September 2020 (feature on corporate buybacks and leverage).

Beach, S., Gamber, W., & Moran, P. (2025). Wealth heterogeneity and consumer spending. FEDS Notes, Board of Governors of the Federal Reserve System, August 5.

Board of Governors of the Federal Reserve System (2017). Equity issuance and retirement by nonfinancial corporations. FEDS Notes, June 16.

Chinco, A., & Sammon, M. (2024). The passive ownership share is double what you think it is. Journal of Financial Economics, 157, 103860.

Chodorow-Reich, G., Nenov, P., & Simsek, A. (2021). Stock market wealth and the real economy. American Economic Review, 111(5), 1613–1657.

DeAngelo, H., DeAngelo, L., & Stulz, R. M. (2010). Seasoned equity offerings, market timing, and the corporate lifecycle. Journal of Financial Economics, 95(3), 275–295.

Graham, J. R., & Harvey, C. R. (2001). The theory and practice of corporate finance: Evidence from the field. Journal of Financial Economics, 60(2–3), 187–243.

Greenwood, R., & Sammon, M. (2025). The disappearing index effect. Journal of Finance, 80(2), 657–698.

Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. American Economic Review, 70(3), 393–408.

Haddad, V., Huebner, P., & Loualiche, E. (2025). How competitive is the stock market? Theory, evidence from portfolios, and implications for the rise of passive investing. American Economic Review, 115(3), 975–1018.

Hawley, J. P., & Williams, A. T. (2000). The Rise of Fiduciary Capitalism: How Institutional Investors Can Make Corporate America More Democratic. University of Pennsylvania Press.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. Econometrica, 50(1), 213–224.

Internal Revenue Code § 4501 (Inflation Reduction Act of 2022, Pub. L. 117-169). One percent excise on net stock repurchases by publicly traded corporations.

Palladino, L. (2019). Ending shareholder primacy in corporate governance. Roosevelt Institute.

Patel, N., & Welch, I. (2017). Extended stock returns in response to S&P 500 index changes. Review of Asset Pricing Studies, 7(2), 172–208.

Poterba, J. M. (2000). Stock market wealth and consumption. Journal of Economic Perspectives, 14(2), 99–118.

Tobin, J. (1969). A general equilibrium approach to monetary theory. Journal of Money, Credit and Banking, 1(1), 15–29.

U.S. Securities and Exchange Commission (1982). Rule 10b-18, 17 C.F.R. § 240.10b-18 (safe harbor for issuer repurchases).

U.S. Securities and Exchange Commission (2020). Response to Congress: Negative Net Equity Issuance.