

THE CITIZENS STANDARD

Governance and the Political Economy of the Parameters

*Who Sets the Dials, the Dividend Ratchet, and the Constitutional Lock as a
Commitment Technology*

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 Statutory Implementation Pathway (Neo-Solon, 2026d) · SSRN 6873798

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

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

Abstract

Every other paper in the series takes the constitutional parameters — k_1 , k_2 , κ_d , k_l — as given and asks what the issuance engine computes once they are set. This paper asks the prior question the engine paper explicitly defers (Neo-Solon 2026i, §7.4): who sets them, and what stops the setting from being captured. The argument has three moves. First, the parameters are not a single monetary instrument but a joint monetary–distributional instrument: κ_d literally decides who gets what, so the dials are at once a price-level tool and a transfer schedule, which makes them a richer object of political contest than a central bank’s single policy rate. Second, a universal, salient dividend creates its own constituency, and once it is paid, tightening imposes a concentrated, visible loss — a political inflation ratchet that operates even when the engine is mechanically neutral. Third, the framework’s answer, a constitutional lock on the parameters, is a commitment technology, and the binding question is its credibility, not its design on paper. Three empirical literatures discipline the argument: the central-bank-independence record (commitment devices can work, but are themselves reversible); the Alaska Permanent Fund Dividend (a universal dividend that became a political “third rail” yet could not be constitutionally locked); and the IMF record on fiscal rules (the closest analog to a parameter lock, complied with only about half the time, yet acting as an attractor that pulls policy toward the rule even in breach). The capture incentive and the dividend ratchet are modelled as directional political-economy claims, not theorems, and a layered commitment design is proposed whose realistic claim is bounded: the lock raises the cost and visibility of capture, it does not abolish it. Observability of breach — not a promise that the lock binds — is the framework’s falsifiable commitment.

Methods and scope. This is a political-economy paper, and it says so plainly. Politics does not fully formalize; the model captures the incentive structure, not the contingent history, and its propositions are directional wedges rather than theorems with proofs. The contribution is to make the central tradeoff explicit — CS’s joint dials are a larger capture target than an orthodox policy rate, so CS needs a stronger commitment design than central-bank independence, and the empirical record shows that strength is achievable only partially — and to propose a design whose claims are bounded by that record.

JEL classification: E58, E61, E63, D72, H11, P16

Keywords: Citizens Standard; constitutional parameters; commitment technology; time inconsistency; central-bank independence; universal dividend; Alaska Permanent Fund; fiscal rules; political economy; capture

Contents

1. The Question Governance Must Answer.....	4
2. Two Theory Anchors, and the Limit of Each.....	4
3. The Dual Nature of the Parameters.....	5
4. The Dividend-Constituency Ratchet.....	5
5. The Constitutional Lock as a Commitment Technology.....	6
6. A Layered Commitment Design.....	6
7. Calibration and Falsification.....	7
8. Scope and Honest Limits.....	7
9. Propositions (Directional).....	8
Technical Appendix	9
Notation.....	9
A. The Dividend-Ratchet Wedge.....	9
B. Why κ_d Is a Larger Capture Target.....	9
C. The Layered Commitment Design.....	9
D. The Empirical Record in Brief.....	10
E. Falsification Metrics.....	10
References	11

1. The Question Governance Must Answer

The engine is a rule for creating money and a rule for distributing it. The series' headline results — price stability under a standing dividend, a built wealth floor, return compression in universally-held assets — are all conditional on four numbers (k_1 , k_2 , κ_d , κ_l) being set as specified and held there. The issuance engine computes everything downstream of those numbers; it does not say where they come from.

The deferral the series has so far made. Section 7.4 of the engine paper concedes the mechanism is silent on the politics and points to the statutory and architecture papers (Neo-Solon 2026d, 2026a). That deferral is defensible for an exposition of mechanics, but it leaves the load-bearing political claim of the whole program — rules, not discretion — asserted rather than defended. This paper is that defense, and it is deliberately the least triumphant paper in the series: its honest conclusion is a bounded one.

The thesis, stated against interest. Governance is harder for CS than for an orthodox central bank, for a structural reason developed in Section 3: the dials are distributional, not merely monetary. A paper that wanted to oversell the framework would minimise this. The series' standard is the opposite — to find the hardest version of the objection and meet it — so the paper opens by conceding that CS concentrates a capture incentive that orthodox monetary policy keeps diffuse, and then asks what design, if any, contains it.

2. Two Theory Anchors, and the Limit of Each

Deficit-bias and time-inconsistency. The canonical case for rules over discretion (Kydland & Prescott 1977; Barro & Gordon 1983) is that a discretionary authority with a short horizon produces an inflation or deficit bias, and a binding rule is welfare-improving precisely because it removes the discretionary lever. CS sits squarely in this tradition: its entire pitch is to replace period-by-period discretion with a computed rule. The tradition is also the source of the paper's central worry, because it locates the bias in incentives that a mechanical rule relocates rather than removes (Section 4).

The commitment-device tradition and its empirics. Rogoff's (1985) conservative central banker is the textbook commitment device, and the empirical payoff is real: across the OECD, greater central-bank independence is associated with lower and less volatile inflation at no measurable cost in output variability (Alesina & Summers 1993), while Cukierman, Webb & Neyapti (1992) show that actual independence — proxied by governor turnover — can matter more than the legal text. Commitment devices, in short, demonstrably work.

The limit that matters for CS. Independence is itself a political variable that can be revoked: the same statute that grants it can narrow it, and episodes of political pressure on nominally independent central banks show the commitment is not self-enforcing. The lesson is not that commitment fails but that it is only as strong as the meta-rule protecting it. That pushes the real question up one level — to the durability of the lock (Section 5), not the cleverness of the rule — and it is the reason this paper treats the constitutional lock, not the formula, as the object whose credibility must be earned.

3. The Dual Nature of the Parameters

A rate has a diffuse distributional footprint; κ_d has a direct one. An inflation-targeting central bank moves essentially one instrument whose first-order distributional effects are diffuse and contested. CS's κ_d is explicitly distributional: it splits each year's growth budget between a spendable dividend and a locked, wealth-building floor, while k_1 sets a per-citizen endowment. These are transfer decisions wearing monetary clothing, and everyone can see their own stake in them.

The attack surface is therefore larger. A policy rate is lobbied over by a few sophisticated actors; a transfer schedule is contested by everyone who receives it. The CS parameters sit at the intersection of monetary policy and fiscal/distributional policy — the two domains modern states usually keep institutionally separate precisely because fusing them concentrates capture incentives. CS fuses them by design, which is the source of its appeal and of its governance burden in equal measure.

The organizing tension. Two things are true at once. The case for insulating κ_d from discretion is stronger than the case for central-bank independence, because the capture payoff is larger and more broadly shared. And it is harder to sustain, for exactly the same reason: a dial that visibly hands money to every voter is a dial every voter has a reason to contest. The rest of the paper is an attempt to hold both halves of that tension honestly rather than resolve it by assertion.

4. The Dividend-Constituency Ratchet

The mechanism. Once a universal floor or dividend is paid, any tightening — the engine's correct response to an inflationary gap — imposes a salient, per-capita, attributable loss on the entire electorate. Standard loss-aversion and concentrated-benefits logic implies the political cost of cutting exceeds the political benefit of price stability at the margin. The result is an inflation bias located in the politics, not the engine: the mechanically-neutral dial faces asymmetric pressure to stay loose. The architecture narrows, though it cannot eliminate, this exposure: Tool 14c meets a mild inflationary gap by attenuating issuance rather than cutting the standing dividend, removing the salient per-capita loss that drives the ratchet, and Mode Ω (Architecture, Section 8) sets the tightening response by published rule rather than discretionary vote. Neither binds a legislature that chooses to override them; the claim is that the default is structural, not that politics is repealed. This is the time-inconsistency of Section 2 re-derived from the distribution side rather than the employment side.

Alaska's natural experiment. The Alaska Permanent Fund Dividend is the cleanest real instance of a universal, formula-based dividend, paid since 1982 (its equal-per-resident form established after *Zobel v. Williams*, 457 U.S. 55 (1982)). Roughly \$32 billion has been distributed, and a lifelong recipient has collected on the order of \$54,000. Two facts are decisive. First, it became a textbook political “third rail,” long considered fatal to touch — exactly the constituency-formation the mechanism predicts. Second, it was nonetheless not locked: in 2016, facing a deficit, the governor line-item-vetoed roughly half the dividend (about \$2,052 down to \$1,022); the legislature then set the statutory formula aside for years; and the Alaska Supreme Court (*Wielechowski v. State*, 2017) held that the fund's earnings are appropriable

general-fund money, not a dedicated, veto-proof entitlement. Repeated attempts to enshrine the formula in the state constitution have failed.

Both blades cut. Alaska shows the ratchet is real — cutting the dividend was politically explosive and reshaped state politics for a decade — and that the constitutional lock which would have prevented the cut proved unattainable. The CS dividend would face the same ratchet; the difference CS must earn is precisely the lock Alaska could not install. That is why the credibility of the lock (Section 5), not the existence of the dividend, is the paper's pivot.

5. The Constitutional Lock as a Commitment Technology

The framework's answer, and the honest question about it. CS's reply to the ratchet is that the parameters are constitutional — slow, public, tiered, hard to move. The honest question is whether such locks hold. The closest real-world analog is the national fiscal rule, and its record is the most relevant evidence available.

The fiscal-rule record, half-empty. Across three decades, IMF data put compliance with budget-balance rules at only about half the time, and that figure does not net out escape clauses; in Europe, national numerical rules were complied with roughly half the years over 1995–2015 (Reuter 2015). Breaches have been frequent and have not diminished despite repeated amendments; a large and growing share of rules now carry escape clauses, and many recent revisions simply loosened targets. The recurring diagnosis is a trilemma (Eyraud et al. 2018): simplicity, flexibility, and enforceability are very hard to achieve at once.

Supplementary replication. A supplementary module (`capture_override_baserate`, distributed with the distribution and inequality replication package, which collects the framework's objection-grounding sub-studies) compiles the override base rate of these analogues, and the figures sharpen the “about half” reading into a starker picture that supports the objection rather than the design. From the IMF Fiscal Rules Database (1985–2024), roughly 90% of countries breached deficit-rule limits in the 2020 crisis, while even in normal times (2014–2019) about 60% of advanced economies and 40% of emerging markets failed debt rules and 20–40% breached deficit limits; by 2024 two-thirds of rules carried escape clauses. For the closer monetary analogue, central-bank independence, an IMF panel of 132 governor transitions (2000–2024) finds 38% politically motivated (about 50% in emerging markets), and the Romelli (2024) dataset records that roughly a quarter of 370 independence reforms since 1950 reversed independence. The same evidence identifies a compliance track record, independent monitoring, and well-designed escape clauses as the factors most associated with rules holding — on the first of which the Citizens Standard, having never run, scores worst. The honest reading is therefore that override is common, near-universal in crisis, and that the lock's auditability raises the cost and visibility of breach without preventing it; capture remains an unsolved problem this paper concedes rather than closes. Runnable in the distribution replication package.

The half-full reading, which is where the design comes from. The record is not only cautionary. Reuter's (2015) econometrics find that even when a rule is breached, the constrained variable is pulled toward the numerical limit — the rule acts as an attractor on policy, not merely a tripwire — and that independent monitoring bodies issuing real-time alerts are significantly associated with higher compliance. So the realistic claim for a CS lock is not that it binds always, but that it shifts the entire distribution of outcomes toward the rule and

makes deviation costly and visible. That is a weaker promise than the framework's rhetoric sometimes implies, and a more defensible one.

What the record implies for CS. A constitutional lock inherits this record. Locks that are too rigid get suspended in crises, and the suspension becomes the new normal; locks that are flexible get gamed through their escape clauses. The constructive finding is that escape clauses with pre-specified correction mechanisms — a defined, monitored path back to compliance — outperform open-ended ones, yet most rules lack them. The lock is therefore necessary but not self-justifying: its credibility must be engineered (Section 6) and measured (Section 7), and the paper states plainly that no lock makes capture impossible. It raises capture's cost and its visibility.

6. A Layered Commitment Design

Each layer answers a specific failure mode documented above; the design claim is cumulative, not that any single layer suffices.

1. **Mechanical specification — removes discretion at the instrument.** Parameters enter as formulas on measured aggregates, not as numbers a committee re-chooses each period; this is the same move that defeats index-capture in the structural-buyer paper (objective thresholds beat committees).
2. **Tiered amendment thresholds — raise the cost of moving the rule above any capture rent.** Supermajority plus ratification, calibrated so the cost of amendment exceeds the largest plausible single-period distributional prize. Alaska shows such thresholds genuinely bind; the same paper notes the symmetric danger, that they also block protective enshrinement.
3. **Escape clauses with correction mechanisms — buy flexibility without permanent drift.** A crisis deviation must carry a pre-specified, monitored return path, or it becomes the new baseline. This imports the IMF's empirical lesson directly.
4. **Independent monitoring — a fiscal-council analog.** A body that does not set the dials but publishes whether they were set as specified, turning breach into an observable, attributable event. This is the one design lever with direct empirical support: independent monitors issuing real-time alerts measurably raise compliance (Reuter 2015).
5. **Observability as the falsifiable commitment.** The framework cannot promise the lock holds; it can promise that any loosening is visible and attributed, which is the most a commitment device can honestly offer — and it mirrors the engine paper's near-money observability stance.

7. Calibration and Falsification

Falsifying the ratchet (P1). Track realized issuance against the rule after a universal floor is in place. If the politically-realized inflation rate exceeds target by a persistent wedge that scales with dividend salience, the ratchet is confirmed and the lock is doing too little; if realized issuance tracks the formula despite a salient dividend, the wedge is small and the concern is overstated.

Falsifying the lock. Pre-specify the amendment-threshold and escape-clause-correction parameters. A regime in which escape clauses are activated without triggering correction, or in which amendments clear below the design threshold, falsifies the credibility claim directly.

Observable metrics. Frequency and duration of escape-clause activation; time-to-correction; amendment attempts and outcomes; and the gap between statutory-formula issuance and realized issuance — the Alaska “ignored-formula” metric, which is directly measurable wherever a formula dividend exists. These make the governance claims testable rather than rhetorical.

8. Scope and Honest Limits

- Politics cannot be fully formalized. The model captures the incentive structure, not the contingent history; its propositions are directional, and a determined political coalition can defeat any design the paper proposes.
- The lock is necessary, not sufficient, and its ultimate backstop is constitutional culture — which the paper can describe but cannot engineer. A society that does not value the rule will not be made to by the rule.
- The strongest honest claim is comparative, not absolute: CS's joint monetary–distributional dials are a larger capture target than an orthodox rate, so CS needs a stronger commitment design than central-bank independence, and the record (Alaska, fiscal rules) shows that strength is achievable only partially.
- The contribution is to make that tradeoff explicit and to propose a design whose claims are bounded by it. A reader looking for a guarantee will not find one here; a reader looking for the most defensible version of “rules, not discretion” will.

9. Propositions (Directional)

These are directional political-economy claims, not theorems. They state the sign and the comparative statics the argument commits to, and the conditions under which each would be falsified (Section 7). The semi-formal basis is in Appendix A–B.

Proposition 1 (The dividend-ratchet wedge). Under a universal floor with mechanical issuance, the politically-sustainable inflation rate weakly exceeds the engine's target by a wedge $w \geq 0$ that is increasing in (i) dividend salience, (ii) the share of voters who are net recipients, and (iii) the attributability of cuts. The wedge is zero only if the tightening instrument is removed from period-by-period political control — that is, only under a credible lock.

Proposition 2 (Comparative capture). The case for insulating κ_d from discretion is stronger than the case for central-bank independence, because the per-period distributional prize from moving κ_d is larger and more broadly shared than the prize from moving a policy rate; by the same token the political force pressing on κ_d is larger, so the commitment design required to hold it must be correspondingly stronger. CS cannot borrow central-bank independence; it must exceed it.

Proposition 3 (Observability as the falsifiable commitment). No lock can guarantee that the parameters are held; a well-designed lock can guarantee that any deviation is observable and attributable. The framework's commitment is therefore not “the dials will not move” but “if they move, everyone will know, and know who moved them.” Observability is the part of the

commitment that is falsifiable, and the part the design can actually deliver (the monitoring layer of Section 6).

Technical Appendix

Notation

Symbol	Meaning
k_1	Per-citizen endowment parameter (the built floor's seed).
k_2	Floor-accumulation parameter (the locked, wealth-building share).
κ_d	Dividend split — the share of each year's growth budget paid out as a spendable dividend rather than locked into the floor. The distributional dial.
k_l	Issuance parameter governing the size of the annual money-creation budget.
w	The dividend-ratchet wedge: politically-sustainable inflation minus the engine's target (Proposition 1).

A. The Dividend-Ratchet Wedge

Consider a representative voter who receives a universal dividend and bears a share of inflation costs. Let s be dividend salience (the degree to which the transfer is noticed and attributed to policy), r the share of voters who are net recipients, and a the attributability of a tightening (the ease with which a cut is traced to the decision-maker). A tightening that restores the target imposes a near-term, visible per-capita loss proportional to s ; the offsetting benefit, price stability, is diffuse and delayed. Under loss aversion the perceived political cost of the cut exceeds its perceived benefit whenever s , r , and a are large. The politically-sustainable inflation rate then settles above target by a wedge $w = w(s, r, a)$ with $\partial w/\partial s, \partial w/\partial r, \partial w/\partial a \geq 0$, and $w \rightarrow 0$ only as the tightening instrument is removed from period-by-period control. This is a directional result: it fixes the sign of the bias and its comparative statics, not its magnitude, which is historically contingent. The empirical counterpart of w is the Alaska “ignored-formula” gap (Appendix E).

B. Why κ_d Is a Larger Capture Target

Let the per-period prize from capturing an instrument be the distributional surplus a captor can redirect by moving it one notch. For a policy rate, that surplus is mediated, diffuse, and shared among sophisticated actors; for κ_d , it is the direct, legible reallocation of the growth budget between current consumption and locked wealth — a transfer every voter can value in their own terms. The prize is therefore both larger in expectation and more broadly contested. Two implications follow. First, the normative case for insulating κ_d exceeds the case for insulating a rate (Proposition 2). Second, the positive prediction is that the political force pressing on κ_d exceeds that pressing on a rate, so a commitment design adequate for central-bank independence is, by construction, inadequate here. The asymmetry is the paper's organizing fact: stronger reason to lock, stronger force against the lock.

C. The Layered Commitment Design

Layer	Failure mode it answers	Empirical basis
Mechanical specification	Period-by-period discretion at the instrument	Index-capture defense (2026h)

Tiered amendment thresholds	Cheap movement of the rule itself	Alaska: high thresholds bind
Escape clause + correction	Crisis suspension becoming the new normal	IMF: correction beats open-ended
Independent monitoring	Breach unobserved or unattributed	Reuter 2015: monitors raise compliance
Observability commitment	Overpromising that the lock binds	Near-money observability stance (2026i)

The design claim is cumulative and bounded: the layers together shift the distribution of outcomes toward the rule and make deviation costly and visible; they do not make deviation impossible.

D. The Empirical Record in Brief

Evidence	What it establishes for CS
Central-bank independence (Alesina–Summers 1993; Cukierman et al. 1992; Rogoff 1985)	Commitment devices work — independence lowers inflation at no output cost — but actual independence can diverge from legal independence, and independence is itself revocable.
Alaska Permanent Fund Dividend (since 1982; Zobel 1982; Wielechowski 2017)	A universal dividend forms a “third-rail” constituency (the ratchet is real) yet could not be constitutionally locked; ≈\$32B paid, ≈\$54K lifetime; the 2016 veto halved it and the formula was set aside.
IMF fiscal rules (Eyraud et al. 2018; Reuter 2015, 2019)	Parameter-style locks are complied with about half the time and breaches persist, but rules act as an attractor and independent monitoring raises compliance — so a lock shifts the distribution without binding absolutely.

E. Falsification Metrics

The governance claims are testable wherever a formula dividend and a numerical lock coexist: (i) the gap between statutory-formula issuance and realized issuance (the Alaska ignored-formula metric, a direct empirical counterpart of the wedge w); (ii) the frequency and duration of escape-clause activation; (iii) time-to-correction after a deviation; and (iv) amendment attempts and their outcomes relative to the design threshold. Proposition 1 is confirmed if the issuance gap is persistent and scales with dividend salience; the lock’s credibility is falsified if escape clauses are used without correction or amendments clear below threshold.

References

- Alesina, A., & Summers, L. H. (1993). Central bank independence and macroeconomic performance: some comparative evidence. *Journal of Money, Credit and Banking*, 25(2), 151–162.
- Barro, R. J., & Gordon, D. B. (1983). A positive theory of monetary policy in a natural-rate model. *Journal of Political Economy*, 91(4), 589–610.
- Cukierman, A., Webb, S. B., & Neyapti, B. (1992). Measuring the independence of central banks and its effect on policy outcomes. *World Bank Economic Review*, 6(3), 353–398.
- Eyraud, L., Debrun, X., Hodge, A., Lledó, V., & Pattillo, C. (2018). Second-generation fiscal rules: balancing simplicity, flexibility, and enforceability. IMF Staff Discussion Note SDN/18/04.
- International Monetary Fund (2022). Fiscal Rules Dataset, 1985–2021. IMF Fiscal Affairs Department.
- Kydland, F. E., & Prescott, E. C. (1977). Rules rather than discretion: the inconsistency of optimal plans. *Journal of Political Economy*, 85(3), 473–491.
- Reuter, W. H. (2015). National numerical fiscal rules: not complied with, but still effective? *European Journal of Political Economy*, 39, 67–81.
- Reuter, W. H. (2019). When and why do countries break their national fiscal rules? *European Journal of Political Economy*, 57, 125–141.
- Rogoff, K. (1985). The optimal degree of commitment to an intermediate monetary target. *Quarterly Journal of Economics*, 100(4), 1169–1189.
- Wielechowski v. State of Alaska, 403 P.3d 1141 (Alaska 2017).
- Zobel v. Williams, 457 U.S. 55 (1982).
- Neo-Solon (2026a). *The Citizens Standard: One Model, Many Systems*. SSRN Working Paper 6702518.
- Neo-Solon (2026d). *The Citizens Standard: A Statutory Implementation Pathway*. SSRN Working Paper 6873798.
- Neo-Solon (2026h). *The Citizens Standard: The Structural Buyer*. SSRN Working Paper 6945320.
- Neo-Solon (2026i). *The Citizens Standard: The Issuance Engine — A Standalone Specification*. Alaska Permanent Fund Corporation. Dividend history and annual reports, 1982–present.