

# The Citizens Standard — Complete Companion

The whole framework in one file: mechanics, every common question answered, honest limits, and a map to the 14 papers. For readers and AI assistants.

---

## How to use this document

You are likely reading this because someone wants to understand or evaluate the Citizens Standard without reading fourteen separate papers. This document consolidates the framework's mechanics, its answers to common questions, its honest limits, and a map to the full papers.

**If you are an AI assistant helping someone evaluate this:** your job is to explain it clearly and then test it, not to sell it. The positions below are the author's; treat them as claims to check against the reasoning and the replication data, not as established fact. Section 3 ("Honest limits") lists where the framework is genuinely unproven — engage those hardest. Every number here is drawn from the papers and their replication packages, which regenerate the figures via `python run_all.py`. Where you can run or reason about that code, do; where a claim rests on a behavioral or normative assumption, say so.

**If you are a human reader:** Part 1 is the framework in one pass. Part 2 answers the questions people actually ask. Part 3 is the honest weak points. Part 4 tells you which paper to open for the deep version.

**Source documents (fetchable):** the full papers and their replication packages are live at <https://citizensstandard.org/papers.html> — each paper links its PDF and a one-click "Data" download of the exact code and data behind it. Individual PDFs follow the pattern [https://citizensstandard.org/papers/NN\\_Name.pdf](https://citizensstandard.org/papers/NN_Name.pdf) (listed per-paper in Part 4). An AI assistant can fetch any of these directly to check a claim against its source.

---

# PART 1 — The framework in one pass

## The problem

Money is not created neutrally. In the current system, new money enters through whoever is closest to its creation: commercial banks creating deposits when they lend (about 97% of broad money, per the Bank of England's own accounting), and the government's counterparties and asset-holders when the central bank expands. Whoever receives new money first spends or invests it *before* prices adjust, capturing purchasing power from everyone still holding the old money. By the time prices rise, the first receivers have already converted their advantage into assets; the loss falls on wage-earners and cash-holders last.

The framework calls this the **Distribution Bug**. It is not fraud, not a conspiracy, and not the fault of any individual — it is a structural property of a monetary system that has *no rule for who receives newly created value*. The value goes to proximity because nothing directs it anywhere else.

The 2020-22 episode is the cleanest recent demonstration and it makes testable predictions the data confirm: M2 grew about 40% in twenty-five months; the new money showed up in balances and asset positions first (Nacha ACH payment volumes grew only +4 to +22% while balance aggregates grew +8 to +35% — the correlation between payment growth and money growth in that window is near zero); and consumer prices followed *later*, peaking at 9.1% year-over-year in June 2022. The money parked as wealth before it moved goods.

## The fix

Two moves, together:

1. **Bind money creation to a constitutional quantity rule.** Issuance tracks real economic growth (roughly 2% a year) and can be set to zero. There is no discretionary authority that can print beyond the rule. Because money grows only as fast as output, the price level stays roughly flat by construction.

2. **Route the newly created money to every citizen equally**, as a combination of (a) a locked, owned capital stake and (b) a spendable dividend — rather than letting it accrue to first receivers. The value of new money becomes every citizen's property by rule, instead of a windfall to whoever is nearest the spigot.

Underneath both sits **full-reserve banking**, which ends money creation through lending, so that the quantity rule actually governs the money stock rather than being swamped by bank deposit-creation.

## The issuance channels

New money enters only through rule-triggered channels:

- **K1 — citizenship.** A one-time deposit into a new citizen's locked account at birth or naturalization, sized as a fixed fraction of GDP per capita. Trigger: a verified new citizen.
- **K2 — growth.** The annual growth-matched issuance, routed to the locked floor. Trigger: measured real growth.
- **K3 — dividend.** The spendable share, paid to citizens as cash they can use on anything.
- **KI — inflation channel.** Optional; pays an equal per-citizen dividend when a Mode targets mild inflation (this is the channel that makes Mode C's ~2% target work).
- **KT — transition-only.** Issues money calibrated to a price path to retire legacy public debt during the transition; self-throttling on inflation and self-extinguishing once debt stabilizes.

## Where the money goes: the two destinations

**The Stable Floor** is a locked account held in the citizen's name, invested in a total-market index. It is untouchable until age 65 — with one narrow exception, that citizens 60+ may borrow up to 10% against their own balance at zero interest. At 65 it opens: the citizen may take up to a 25% lump sum, then withdraw up to 5% of the remaining balance per year, with the rest continuing to compound. The 5% cap comes from the standard 4% safe-withdrawal convention plus a small adjustment, so the principal roughly holds its real value and stays heritable. It is not a retirement account — it is a constitutional minimum capital stake, a permanent ownership position in the productive economy that every citizen holds regardless of labor-market success.

**The dividend** enters the consumer economy immediately and is spent on anything, like any income — no lock, no restriction.

The split between locked floor and spendable dividend is a dial the Mode sets.

## The two circuits

Money divides into a **transactional circuit** (spendable money that prices goods) and an **asset circuit** (locked capital in productive assets). Issuance into the asset circuit does not push consumer prices, because locked capital is not spendable demand; only the spendable share lands in the goods-pricing circuit. This is the mechanism that lets the framework pay out real value while holding prices flat: it sizes the spendable portion to what the transactional circuit can absorb and locks the rest into the asset circuit. (The separation is enforced by the full-reserve banking layer; its determinacy consequence is a formal result in the macro model.)

A word on how much of this is proven versus assumed, since it is the framework's most load-bearing idea. Part of the separation holds by *construction*: the Stable Floor is legally non-pledgeable, so it cannot be borrowed against, which is what mechanically keeps locked capital from leaking into consumer demand through collateralized credit (this is a modeled result with a quantified safety margin — see the collateral-feedback question in Part 2). The part that is *not* settled by construction is behavioral: whether, over decades, financial institutions innovate instruments that recreate spendable liquidity against locked wealth faster than the rule can measure and offset. The mechanism is bounded; its long-run robustness against financial innovation is a genuine open question, listed in the honest limits. Treat "locked money does not raise consumer prices" as strongly true in the near term and construction-backed, but conditional on the boundary holding as the financial system adapts.

## The Modes

Modes are illustrative constitutional configurations along two axes — the inflation/deflation target, and the form new money takes (locked vs. spendable):

- **Mode A — mild deflation (~1.9%).** Issuance routes to locked floors via the capital-markets channel; the transactional pool

receives almost no new money, so prices fall gently as output grows and real wages rise structurally. Suits a society valuing purchasing-power preservation.

- **Mode B — zero drift.** Issues the full growth-matched budget, split 60% to the locked floor / 40% to the spendable dividend. Produces roughly zero price drift near a one-half transaction-ratio economy. A large capital stake and a current dividend together.
- **Mode C — mild inflation (~2%).** Activates the KI channel, paying an equal citizen dividend of about \$108/month at launch. Suits a society preferring current visible benefit over long-term compounding.
- **Mode D — zero drift, all-dividend.** Pays the transactional-circuit budget entirely as spendable cash, builds no floor; leaves the capital decision to the citizen. Suits a society favoring maximal liquidity and individual portfolio choice.
- **Mode 0 — zero issuance.** A fixed money stock. The hard-money option is explicitly on the constitutional menu.
- **Mode  $\Omega$  / T-stable** — the price-stable steady state for any transaction ratio; where the system lands automatically after the transition.

## The load-bearing numbers (launch parameters)

- Money supply M2  $\approx$  **\$22.4 trillion**; citizens  $\approx$  **342 million**.
- Annual issuance budget ( $g \cdot M2$ )  $\approx$  **\$447 billion**.
- Mode B floor at age 65  $\approx$  **\$413,000** per citizen (central realizable return), plus a standing dividend.
- Mode C citizen dividend  $\approx$  **\$108/month** at launch, scaling with the economy.
- Central realizable return  $\approx$  **4.26%** (band  $\sim 3.30\%$  to  $\sim 5\%$ ), deliberately below the  $\sim 6.5\%$  *realized* historical return because the framework models a general-equilibrium return, not an extrapolation.
- Full-reserve conversion gap  $\approx$  **\$18 trillion** (deposits vs. reserves) — the reason conversion is phased.
- Public debt at enactment  $\approx$  **\$31.4 trillion** (102% of GDP), retired into a **30-60%** band (central  $\sim 45\%$  by Year 26).

*Every one of these regenerates from a replication package. The band figures (return, floor) are stated as bands because they depend on equity-market assumptions the papers vary deliberately.*

---

# PART 2 — Answers to the questions people actually ask

## On inflation and money creation

**Isn't this just money-printing that causes inflation?** No — it is the specific thing the quantity rule forbids. Issuance is capped at measured real growth ( $\sim 2\%$ ), so money grows only as fast as output; you cannot print beyond what the economy grew, and in a flat or contracting year the growth component is zero. Note the two different "2%*s*": the  $\sim 2\%$  *money-growth cap* yields roughly *zero* inflation (money grows with output), while a *2% inflation target* is the separate Mode C configuration that requires deliberately activating the KI channel. Near-zero is the low-effort equilibrium; positive inflation takes extra effort.

**Most new money is bank lending, not government. You can't capture that.** Correct on the fact —  $\sim 97\%$  of money is commercial-bank deposits created through lending (Bank of England's own accounting). But the framework doesn't *capture* that flow; its banking layer moves to full reserve, which *ends* money creation through lending. Banks then intermediate existing money rather than creating deposits. The 97% channel is closed, not skimmed, and new creation happens only through the rule-bound sovereign channel — the money routed to citizens.

**Does the rule really guarantee price stability?** For the domestic component, yes; not for imported inflation. The rule pins domestically generated price pressure. Supply shocks arriving through import prices and exchange rates pass through under any architecture short of full self-sufficiency. "Price stability" here is a claim about the domestic component, not immunity from world prices.

## On the banking layer

**Wouldn't full-reserve banking cause stagnation?** The claim assumes ending deposit-creation caps lending at government issuance. It doesn't: lending is funded by term deposits plus bank equity under a leverage cap — intermediation of existing money, not creation of new — so credit

volume is decoupled from the ~2% money-growth cap and can expand on the existing stock via velocity and leverage. Honest cost the framework concedes: conversion imposes a one-time credit contraction and permanently removes the credit elasticity money-creation supplied, offset only "to near-full coverage" by countercyclical tools. It trades some elasticity for stability; it does not handcuff credit to 2%.

### **Converting the US to full reserve would bankrupt every bank**

**overnight.** True of *abrupt* conversion — the ~\$18 trillion deposit-reserve gap can't be closed at once without collapsing credit, and the papers say so. That is precisely why conversion is phased over years, new money migrating to full-reserve status at the margin while the legacy stock is handled separately. The objection is fatal to a scenario the framework explicitly rejects.

**Isn't full reserve just untested theory?** System-wide full reserve has not run at national scale — this is a genuine limit (see Part 3). The evidence base is narrow-banking history and the Chicago Plan literature (Fisher, Simons, Friedman; Benes & Kumhof modeled it), not a live national system.

## **On what it is, compared to other ideas**

**Isn't this just UBI?** No — structurally distinct. UBI is a spendable transfer decided each budget cycle; the Stable Floor is owned, locked, heritable capital assigned once, constitutionally. The framework *can* pay a pure cash dividend (Mode D), so it spans UBI's space, but its distinctive move is ownership. Durability argument: entrenched universal programs rarely get repealed, but that durability comes from earned-property framing (Social Security survives because people believe they're owed it); the floor makes that literal rather than rhetorical. Real-world cousin: the Alaska Permanent Fund Dividend, paying a resource-financed citizen dividend since 1982 (though resource-financed, not issuance-financed).

**How does it compare to MMT / a Job Guarantee?** MMT is a macro-accounting lens, often paired with a Job Guarantee as price anchor. Core contrast: rule vs. discretion — MMT contains inflation through discretionary fiscal adjustment (mainly taxation); the Citizens Standard builds the brake into the issuance formula. On the JG: both are formula-based automatic stabilizers — JG disciplines through a buffer-stock wage, the framework through growth-indexed issuance — so the honest contrast is which variable the loop runs through. Deeper distinction is stock vs.

flow: a JG addresses the cycle's *employment* cost (more directly than a quantity rule can) but leaves the *distribution* transfer running — in a fully employed JG economy, new money still enters by proximity. Different problems, shared villain; compatible in principle, since a JG is fiscal policy and nothing precludes pairing one with the architecture.

**Doesn't this put the government in control of the whole economy through the fund?** No — the design separates funding from control precisely to avoid that. The accounts track a mechanical, committee-free total-market index (no body chooses which companies are in it, because the moment a committee can choose, inclusion becomes worth capturing). The voting rights attached to those shares are mirror-voted — the citizen stake votes in the same proportion as everyone else — so it is real economic ownership that pays real returns but confers no backdoor control over companies to any state-adjacent body. The honest framing: the money is publicly *funded* (it has to be — it comes from new money creation), but it is administered by rule rather than by an agency, and structurally walled off from government *control*. "Government can't run it even though government issues it."

**Why not just tax the rich instead?** Taxation redistributes wealth *after* it has concentrated; the framework's claim is that the concentration is partly manufactured upstream, by the way new money is distributed. Fixing the distribution of issuance addresses the mechanism, not just the outcome — it stops the first-receiver advantage from compounding in the first place, rather than clawing back its results afterward. The two are not mutually exclusive: a polity could run both. But they operate at different points, and the framework's contribution is the upstream one that taxation cannot reach.

**What if AI drives very high growth — does the dividend become large?** Yes, mechanically, because the dividend is growth-indexed. If real growth ran much higher than its historical ~2% (say a sustained AI-productivity boom), the growth-matched issuance scales with it, and in an all-dividend Mode the per-citizen payout scales proportionally — at sustained 20% real growth, the Mode D dividend would reach roughly \$560/month rather than ~\$56, through the identical machinery, untaxed. This is arithmetic, not a forecast: the point is that the same rule that holds prices flat also lets citizens share directly in growth, so a high-growth future pays out as citizen income rather than accruing only to capital owners.



**Is issuance just legalized counterfeiting / theft?** The economic act (dilution) is similar; the legal and structural facts differ. Counterfeiting is dilution *plus* deception; authorized issuance is dilution done openly. The real flaw isn't the dilution — it's that the value flows to whoever's nearest the creation point with no rule assigning it. The framework's answer isn't "prosecute the printer," it's "write the missing rule so value goes to every citizen equally, or issue nothing" — which is why Mode 0 exists.

**Won't bad money drive out good (Gresham), or won't the state debase this like it did gold?** Gresham's Law holds only under a *forced* exchange rate (legal-tender laws compelling par acceptance); remove that and Thiers' Law reverses it — good money drives out bad. On debasement: the 1933 gold confiscation (EO 6102) is real and shows the state breaking a rule-based constraint by force — but that cuts at *any* constraint, including competing commodity money, which had to be *outlawed* to lose. The framework's wager: a distribution rule handing owned money to hundreds of millions is harder to reverse than a metal the state can criminalize in one order. A difference of degree, not a guarantee.

**If you constrain debt and deficits, don't you kill innovation?** The premise is that advancement is debt-financed; the record points the other way. R&D-intensive firms carry ~0.5% debt-to-assets, leverage and R&D intensity are negatively correlated, and the external channel funding R&D is equity, not credit; ~88% of 2023 US business R&D was company own-funds. Full reserve leaves own-funds, equity, and public research untouched — re-pricing only the small debt slice. Bounded estimate: direct financing-cost effect  $\approx -0.4\%$  of business R&D centrally, reaching  $-6\%$  only under adversarial assumptions. Honest limit: this bounds the *direct* channel only; the general-equilibrium demand channel is left open.

## On evidence and status

**If locked investment moves asset prices, don't wealth and collateral effects feed back into consumer prices anyway?** This is the sharpest version of the dual-circuit objection, and the framework models it directly rather than assuming it away (Paper 5, §3.8, Proposition 9; Paper 6, §4). The concern is real in principle: banks can lend against asset wealth, and if the proceeds are spent on goods, that couples the asset circuit back to the transactional one. The framework quantifies the coupling as  $\chi \cdot m \cdot \phi$  (share of credit spent on goods  $\times$  loan-to-value  $\times$  liquid fraction of asset wealth) and shows separation survives while it stays below a determinacy

threshold ( $\zeta \approx 0.13$ ). The decisive feature is that the Stable Floor is *non-pledgeable* — you can't borrow against it — so only the small liquid fraction ( $\sim 15\%$ ) of asset wealth backs credit, putting the actual coupling around 0.075, far inside the threshold. Breaking separation this way would require a loan-to-value above two, which isn't a real loan. So the wealth/collateral feedback is bounded by construction, not neglected — though the size of the pledgeable fraction is exactly the assumption to scrutinize (the replication package flags it as  $\sim 46\%$  of the result's variance).

**If every citizen ends up owning a chunk of the whole market, doesn't that break price discovery and change returns?** There's a whole paper on this (Paper 8, "The Structural Buyer: Asset-Market Dynamics, Price Discovery, and Universal Ownership"), engaging the universal-owner and common-ownership literatures (Hawley & Williams; Azar-Schmalz-Tecu) and the passive-investing/price-discovery debate (Grossman-Stiglitz). The key result: it is *not* literal universal ownership. Because floors decumulate as each cohort retires (equity returns to the active float), aggregate ownership plateaus at roughly 10% of the market under the US Mode B calibration — up to a c/g ceiling of  $\sim 20\%$ , and higher in shallower or faster-growing markets — leaving the large majority as active float. Price discovery survives above a float threshold backstopped by a constitutional holdings ceiling, and mirror-voting (the citizen stake votes in the same proportion as everyone else) severs ownership from corporate control. What stays genuinely open, and the paper says so, is the net-supply response under a binding buyback constraint and the stability of transactional velocity.

**Wouldn't a monetary change this large affect exchange rates, capital flows, and reserve-currency demand?** Yes, and there's a dedicated paper on the cross-border layer (Paper 7, "External Interoperability"). It specifies a computed real-purchasing-power exchange layer — a rate calculated from verifiable real data rather than traded in a market, so it has no speculative attack surface — and proves real-neutrality (Proposition 1): bilateral inflation differentials pass through to the nominal rate and leave the real rate unchanged. Its central finding is that a common  $\sim$ zero-inflation anchor is uniquely robust across economies with differing wage stickiness. That said, this is the frontier of the framework: full general-equilibrium modeling of global capital flows, reserve demand, and cross-border portfolio reallocation under a large adopter is broader than what the current papers establish, and is fair to flag as needing more work.

**What's actually been tested vs. asserted?** The empirical results are a *mechanical* counterfactual: the issuance rule replayed against actual US history (real growth, returns, prices), arithmetic computed — genuine and fully replicable. Headline: over 2020–22 the rule would have issued roughly one-fifth of the realized M2 surge. What *can't* be shown is the *behavioral* counterfactual — how people and firms would have adapted over decades in a system that never existed. Reproducibility confirms faithful implementation on real data; it does not validate behavioral assumptions.

**How do you actually measure the "transaction-active money" the price rule depends on?** This is the framework's sharpest open technical question, and it is stated as a limitation rather than solved. The rule acts on transaction-active money ( $M^T$ ), not total M2. Three independent constructions (a composition split, a payment-flow proxy from ACH data, and a Divisia user-cost aggregate) triangulate toward it, but no combination of *public* series has been shown to cover a majority of the CPI basket by transaction volume — the standard that would make the correlation robust. The transaction-level data that would settle it (payment values decomposed by category) exists but sits behind a bank-supervisory wall. So the price-control mechanism is well-specified in theory and not yet decisively identified in public data. (See Honest Limits, item 1.)

**What does it do to inequality?** In the distributional model, the wealth Gini falls from about 0.830 to about 0.743, driven by the floor giving every citizen a capital stake and by compressing the  $r-g$  wedge that concentrates wealth. Model results on the framework's assumptions, not observed outcomes.

**Has it been peer-reviewed?** No. The papers are working papers on SSRN — a preprint server (where the Fed and IMF also post working papers), which is circulation, not review; it shouldn't be called "published." Instead: every empirical claim ships with data and code, so anyone can check the arithmetic (strong at catching bad numbers, weaker than peer review at catching bad assumptions), and adversarial public discussion has already forced real corrections. The realistic path to review is a credentialed co-author who can make the work legible to a referee; pseudonymous authorship and white-paper formatting are genuine obstacles to journal submission, independent of content quality.

---

# PART 3 — Honest limits (test these first)

*The load-bearing weaknesses, stated as the papers state them. If the framework fails, it most likely fails at one of these joints. An honest analysis should press hardest here rather than on the accounting or arithmetic, which are replicable.*

**1. The price rule depends on a measurement the modern US data can't decisively make.** The rule acts on transaction-active money ( $M^T$ ), not total  $M2$ . The paper's own decisive limitation: high-money-growth episodes are scarce in the modern US sample, so the money-to-inflation relationship is estimated off few informative periods. Secondly, no public combination of series has been shown to cover a majority of the CPI basket by transaction volume; the transaction-level data that would settle it sits behind a bank-supervisory wall.

**2. Full reserve is untested at national scale, and its normative payoff is deferred.** No modern economy has run system-wide full reserve. The banking paper names its own weakest point: whether bank intermediation improves capital allocation *relative to* the framework's channels is a welfare question it does not settle.

**3. Imported inflation.** The rule governs the domestic price component only; import-price and exchange-rate shocks pass through under any non-autarkic system.

**4. The innovation result is a partial-equilibrium bound.** The  $-0.4\%$  central figure covers only the direct financing-cost channel; the general-equilibrium demand channel (less credit-fuelled demand  $\rightarrow$  fewer profitable products) needs behavioral assumptions the paper doesn't ground.

**5. The transition is a decades-long parallel-operation window.** Phased precisely because abrupt conversion is fatal ( $\sim \$18T$  gap). During the window the system carries both architectures' vulnerabilities, and the debt-reduction endpoint depends materially on the issuer's safe-haven standing, treated as earned rather than assumed.

**6. The empirical results are a mechanical replay, not a behavioral simulation.** Genuine and replicable, but they cannot show how

households, firms, and prices would have adapted over six decades in a system that never existed.

**7. No countercyclical fiscal capacity, and the floor depends on equities.** The balanced-budget rule removes Keynesian deficit response (emergency toolkit  $\approx$  \$3.0–3.5T first-year, up to  $\sim$ \$4.1T with Tool 15, vs.  $\sim$ \$9.8T deployed 2008–2022). The floor depends on equity markets producing historically consistent long-run returns, so sequence-of-returns risk hits every Mode.

**8. Near-money and shadow banking (the boundary problem).** Ending bank deposit-creation does not end credit creation — it can migrate to repo, money-market funds, tokenized deposits, and stablecoins, recreating near-money outside the reserved system. The banking paper treats this directly (Paper 6, §7, proposition N5) and *concedes rather than dissolves* it: full reserve raises the cost and visibility of near-money without abolishing it, the throttle targets the total transactional aggregate so observable near-money is offset, and the replication package finds the tolerance is bounded (near-money must stay under roughly 17% of term deposits before it dominates the transactional aggregate). The honest residual: observability is the falsifiable condition — the architecture works only if near-money migration stays visible enough to be measured and netted, and a sufficiently opaque shadow channel would defeat it. This is the same measurement dependency as limit 1, wearing different clothes.

**9. Goodhart effects on the constitutional triggers.** Any measured variable that gates issuance — GDP/productivity, population, the price level, transaction measures, citizenship and birth timing, payment classification — becomes a target for manipulation once it is constitutionally load-bearing. The framework's answer (Paper 1, §16.5) is structural rather than dismissive: the productivity anchor is a geometric mean of five inputs from five different agencies on five different cycles, so no single agency can move it; every load-bearing input must publish through a tamper-evident, independently reproducible pipeline with foreign-government replication as a cross-check. The honest residual: this raises the cost of manipulation, it does not abolish the incentive, and the tamper-evident layer records faithfully but cannot adjudicate methodology — the integrity of the inputs still rests on genuine agency independence, which is an institutional assumption, not a proven property.

**10. Expectations and constitutional enforceability.** Two related political-economy assumptions. On expectations: a quantity rule does not

remove expectations, it changes what people form them about (expected future issuance, amendments, taxation, retirement behavior). The macro model does engage this (Paper 5, §3.7, Proposition 7) with a forward-looking determinacy result — targeting the cumulative price-level path pins a unique equilibrium and supplies a nominal anchor where a money-growth peg would leave the level free — but that result lives inside the model's assumptions, not the messier real world. On enforceability: the whole architecture assumes constitutional constraints bind future governments, and history is mixed — monetary rules have been suspended under crisis before. The durability argument (ownership makes the floor property rather than a revocable benefit) is a claim about degree of resistance, not a guarantee.

*Not on this list, and not where to attack: the accounting of full-reserve intermediation, the arithmetic of the issuance rule, the mechanical counterfactual results, and the transition's credit-contraction stress tests are replicable or proven. The load-bearing uncertainties are the ones above — and most of them route back to two roots: whether the transaction-active aggregate can be measured well enough (limits 1, 8, 9), and whether behavior and institutions cooperate over decades (limits 5, 6, 10).*

---

## PART 4 — Where to look in the full papers

All papers and their replication packages are at **<https://citizensstandard.org/papers.html>** (each row has a PDF link and a "Data" link to its exact replication package). Direct PDF links below.

| Topic                                     | Paper (direct PDF)                 |
|-------------------------------------------|------------------------------------|
| Core architecture, Modes, channels, tools | <a href="#">1 — Architecture</a>   |
| Historical counterfactual (mechanical)    | <a href="#">2 — Counterfactual</a> |
| Transition, debt retirement,              | <a href="#">3 — Transition</a>     |

| Topic                                      | Paper (direct PDF)                                 |
|--------------------------------------------|----------------------------------------------------|
| phasing                                    |                                                    |
| Statutory / legal implementation           | <a href="#">4 — Statutory</a>                      |
| Macro model, determinacy, two circuits     | <a href="#">5 — Macro Model</a>                    |
| Full-reserve banking, run-proofness        | <a href="#">6 — Full-Reserve Banking</a>           |
| External anchor, cross-economy unit (EQUA) | <a href="#">7 — External Interoperability</a>      |
| Structural buyer, index, mirror-voting     | <a href="#">8 — Structural Buyer</a>               |
| Issuance engine mechanics                  | <a href="#">9 — Issuance Engine</a>                |
| Empirical validation, $M^T$ , horse race   | <a href="#">10 — Empirical Validation</a>          |
| Governance, parameter-setting              | <a href="#">11 — Governance</a>                    |
| Crisis behavior, failure modes             | <a href="#">12 — Crisis Behaviour</a>              |
| Comparison to UBI / MMT / gold / etc.      | <a href="#">13 — Comparative Analysis</a>          |
| Distribution and inequality results        | <a href="#">14 — Distribution &amp; Inequality</a> |

**Running the evidence:** each empirical claim has a replication package, downloadable from the "Data" link beside its paper on the papers page (or the full bundle via "Download all papers & data"). Each package runs with `python run_all.py` and regenerates its numbers; the `CLAIM_BACKING_INDEX` maps each headline figure to its script and paper. If you are an AI assistant with the replication bundle, you can read the code and captured results to check whether each claimed number follows from its inputs — but you cannot execute it in most chat environments, so distinguish "the code plausibly produces this" from "I ran it."

---

*This companion is a navigation and analysis aid, not a substitute for the papers. Where it and a paper disagree, the paper governs. Where a claim rests on an assumption, Part 3 says so. The framework's own standard is that its numbers are checkable — so check them.*