

THE CITIZENS STANDARD

Full-Reserve Banking and the Two-Circuit System

Payment-Credit Separation, Complete Monetary Control, and Credit Supply under Full Reserve

Neo-Solon

Neo-Solon@hotmail.com

WORKING PAPER — JUNE 2026

COMPANION PAPERS

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

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

The Citizens Standard: 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: External Interoperability and the Common Anchor (Neo-Solon, 2026g) · SSRN 6939600

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: Empirical Validation of the Transactional Aggregate (Neo-Solon, 2026j) · SSRN 6973298

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

The banking architecture of the Citizens Standard is full-reserve. Transaction accounts are held one-for-one against public reserves and are constitutionally protected; term deposits are at-risk investment claims that fund lending; and commercial banks, by construction, cannot create new money by lending (Neo-Solon, 2026a, §9). This paper specifies what that architecture buys and costs in the banking layer of the two-circuit model, and derives five results. First, because banks create no transactional money, the price-relevant aggregate equals sovereign outside money exactly; the inflation-gap throttle (KI; specified in Paper 1, §6.1, and analyzed in Paper 5) controls it directly, and the forward-looking determinacy of the macroeconomic model carries over with no augmented aggregate, no deposit offset, and no saturation failure mode (N1). Second, bank credit is intermediation of pre-existing funds, not money creation; the only channel by which it could re-mix the circuits — credit extended against pledgeable asset wealth and spent on goods — is bounded by the non-pledgeability of the locked floor far below the separation threshold, so circuit separation survives endogenous credit (N2). Third, bank credit equals term deposits plus equity under a leverage cap; full reserve binds only the bank-intermediated slice of credit and leaves market credit untouched; the one-time contraction that full-reserve conversion imposes is offset by the framework's own sovereign issuance through the citizen channels to near-full coverage, and the steady-state cost — a loss of the credit elasticity that money creation supplied, since saving is rate-inelastic — is restored by the countercyclical channels and the leverage rule without restoring money creation (N3). Fourth, run-proofness is structural: fully reserved transaction accounts cannot be run, at-risk term deposits are not demandable money, and the dominant household asset is locked equity outside the run technology, so the maximum systemic money contraction is bounded by the term-deposit share rather than by the whole deposit base (N4). Fifth, the standing critique — the boundary problem — is conceded rather than minimized: 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 observability is the falsifiable condition (N5). The banking reform is the full-reserve tradition of Fisher, Simons, and Benes and Kumhof; the framework's departure is distributional — it routes seigniorage to citizens as locked equity rather than to the state.

JEL classification: E51, E58, G21, E42, E44, G28

Keywords: Citizens Standard; full-reserve banking; Chicago Plan; payment–credit separation; sovereign money; two-circuit monetary system; bank runs; near-money

Calibration
 Saving rate-elasticity
 Empirical literature (~0)
 Calibration
 Credit-demand elasticity
 Empirical literature (~3-4%/pp)
 Calibration
 Pledgeable fraction (0.15)
 Empirical / Papers 1,5
 Calibration
 Separation threshold zeta* (~0.13)
 Paper 5 (2026e, A.12)
 Calibration

Table. Provenance of each quantity used in this paper. Anchors are empirical values from official statistics; structural assumptions are constitutional rules of the framework that hold by construction; calibrations are values set from the literature or the launch split, whose magnitudes (not the qualitative results) depend on the chosen figure. Section 5.4 and the replication package report the loanable-funds frontier under low, baseline, and high elasticity to confirm the result is not an artifact of one calibration. The package further verifies each proposition against a derived two-account balance sheet, distinguishes the structural propositions (which hold as identities) from the calibration-dependent ones, discovers the separation and near-money thresholds algorithmically, and runs a randomized invariance test over ten thousand admissible calibrations: the structural propositions hold in every draw, while the calibration-dependent ones report their pass rate over the admissible space. Beyond verification, the package discovers the model's boundary quantities algorithmically — the maximum term-deposit share preserving an adequate transaction aggregate (0.75 of M2), the maximum credit intensity preserving separation (0.33), the maximum observable near-money before it dominates the transactional aggregate (about 17 percent of term deposits), and the collateral shock separation survives (2.2 times the baseline pledgeable fraction) — and answers parametric what-if queries directly, so the model can be explored beyond the baseline calibration rather than only confirmed at it. A parameter-importance analysis over twenty thousand draws decomposes the variance of each proposition's margin across its assumptions: circuit separation is driven chiefly by the pledgeable fraction (about forty-six percent of the margin variance), while credit capacity is set by the money-stock scale and the term-deposit share rather than the leverage cap, which contributes only about two percent — identifying, for each result, the assumption a reader should scrutinize most.

5.1 The loanable-funds identity and the leverage cap.....	8
5.2 Full reserve binds only the bank slice.....	9
5.3 The conversion contraction and its sovereign offset.....	9
5.4 Steady-state credit elasticity and the countercyclical channels.....	9
6. Payment–credit separation and run-proofness.....	10
7. The boundary problem: near-money.....	11
8. Discussion and relation to the series.....	12
9. Conclusion.....	12
Appendix.....	13
A.1 Notation.....	13
A.2 Propositions.....	14
A.3 Proofs.....	14
A.4 Calibration and falsification.....	15

References..... 17

1. Introduction

The preceding papers of this series describe an architecture for outside money — money issued by the public authority. The Citizens Standard partitions that money by use: a transactional circuit prices goods and services, while an asset circuit holds citizens' equity claims, and the framework's headline results follow from keeping the two circuits structurally distinct. The macroeconomic model (Neo-Solon, 2026e) shows that new issuance landing in the asset circuit raises asset values rather than consumer prices (Proposition 1), that the price level is determinate under a money-quantity anchor (Proposition 7), and that the separation survives a small leak between the circuits (Proposition 9).

A reader who accepts those results is entitled to a sharp question about banks. In the incumbent system, most money is not outside money: commercial banks create it privately. The Bank of England's account is now the orthodox one — whenever a bank makes a loan it simultaneously creates a matching deposit, so bank lending creates deposits rather than intermediating pre-existing ones (McLeay, Radia, and Thomas, 2014; Jakab and Kumhof, 2015) — and roughly four-fifths of the broad money stock in the United States is, by this mechanism, a private bank liability. If that mechanism operated inside the Citizens Standard, the price level would no longer be pinned by the authority's issuance, credit extended against asset-circuit wealth could re-mix the circuits, and demandable deposits would reintroduce the fragility the architecture is meant to remove.

The Citizens Standard answers this question not by bounding private money creation but by abolishing it. Its banking architecture is full-reserve (Neo-Solon, 2026a, §9, §17.3). Transaction accounts are held one-for-one against public reserves; term deposits are at-risk investment claims that fund lending; and banks cannot create new money through lending — they lend only term-deposit liabilities and equity. This is the full-reserve, or 100-percent-money, tradition: Fisher (1936) proposed full reserves against checking deposits to sever money creation from lending, Simons (1936) made it central to the Chicago Plan, Friedman (1948) endorsed it, Tobin (1985) advanced a related deposited-currency proposal, and Benes and Kumhof (2012) revived it analytically, finding in a calibrated model of the US economy that a Chicago-Plan transition could reduce debt and dampen the credit cycle. The framework adopts the load-bearing feature of that tradition and departs from it in one respect only: where the 100-percent-money literature routes the resulting seigniorage to the state, the Citizens Standard routes it to citizens as locked individual equity (Neo-Solon, 2026a, §17.3).

That choice reframes every banking question this paper asks. The original worry — does private money creation defeat the architecture? — does not arise, because there is no private money creation. The questions that remain are the ones a full-reserve system must actually answer, and this paper takes them in turn: whether monetary control is complete when banks cannot create money (Section 3); whether bank credit, now pure intermediation, can still re-mix the circuits (Section 4); how credit is supplied when banks may lend only what savers term out, and at what cost (Section 5); whether the payment system is genuinely run-proof, and on what bound (Section 6); and whether the separation of money from near-money survives the boundary problem that every narrow-banking proposal confronts (Section 7). The paper sits at the intersection of the full-reserve tradition above, the money-creation literature (Tobin, 1963; McLeay, Radia, and Thomas, 2014), the bank-run literature (Diamond and Dybvig, 1983), and the standing critique of narrow banking (Goodhart, 2008). Its contribution is to work those

questions through the specific structure the framework already imposes — a locked, non-pledgeable equity floor as the dominant household asset, and a sovereign issuance engine with its own countercyclical channels.

A note on what is reused. Throughout, the locked floor does double duty. It is the device behind the demographic, labor, and circuit-separation results of the earlier papers; here it reappears as the reason bank credit cannot re-mix the circuits (Section 4) and as a third layer of run-proofness (Section 6). The sovereign issuance engine does the same: the channels that distribute new money to citizens (Neo-Solon, 2026e) are the channels that offset the credit contraction full-reserve conversion imposes (Section 5). The banking results are consequences of structure the framework already has, not new machinery bolted on.

2. Full-reserve banking in the two-circuit system

2.1 Two account types: reserved transaction accounts and at-risk term deposits

The Citizens Standard restructures banking around a single principle: the payment system is a public utility and must never fail, while the credit system is a market activity that may fail in ordinary ways without bringing the payment system down with it (Neo-Solon, 2026a, §9). Two account types replace the current deposit structure. Transaction accounts are full-reserve: balances are held one-for-one against reserves at the issuing authority and are constitutionally protected from creditor claims, bank insolvency, and seizure, so they are functionally identical to physical cash in digital form. Term deposits are at-risk: a citizen may place funds for a fixed term at a competitive return, those funds explicitly fund bank lending and explicitly bear credit risk, and there is no guarantee on them. The depositor knows this and is compensated by the return. The term "full reserve" throughout this paper refers to the money layer: the means of payment — transaction balances — is reserved one-for-one, in the 100-percent-money sense of Fisher (1936) and Simons (1936). Term deposits are deliberately *not* reserved, because their economic function is to be lent; they are investment claims rather than money, and the bank holds reserves only against the transaction accounts. A full-reserve system in this tradition reserves fully against money, not against every liability a bank issues.

The accounting consequence is the one this paper turns on. When a bank lends, it lends term-deposit liabilities and its own equity; it does not credit a new transaction balance into existence. A loan transfers existing reserves from the lender's funding base to the borrower; it does not expand the transactional money stock. In the language of the money-creation literature, the Citizens Standard places banks back into the intermediation-of-loanable-funds mold that the literature shows the incumbent system does not occupy (McLeay, Radia, and Thomas, 2014; Jakab and Kumhof, 2015) — not as a description of how banks behave when left alone, but as a binding statutory constraint. The leverage of the credit system is capped: the protocol sets a maximum of four to one in normal conditions, tightening to three to one when private credit growth outruns nominal output and loosening to five to one when it falls behind (Neo-Solon, 2026a, §9.2).

2.2 The transaction aggregate under full reserve

Write outside transactional money — the authority's issue — as M_o , and asset-circuit money as M_A . In a fractional-reserve system one would add a second layer, the transactional part of

bank deposits, and the price-relevant aggregate would be the sum of the two. Under full reserve that second layer is identically zero: banks create no spendable deposits, so the transactional aggregate is

$$M_T = M_o,$$

and the price level is $P = M_T \cdot V/Y = M_o \cdot V/Y$, set entirely by the authority's issuance relative to output. The contrast with the incumbent system is stark rather than marginal. In early 2026, commercial-bank deposits were the dominant component of the money stock — on the order of \$18 trillion against an M2 near \$22.4 trillion (Federal Reserve releases H.6 and H.8), roughly four-fifths of it a private bank liability. Under the Citizens Standard that share is zero — the transactional money is sovereign reserves, term deposits are investment claims that do not circulate as means of payment, and the dominant household asset, the locked floor, is equity that is neither money, nor a bank liability, nor pledgeable.

2.3 Relation to the Chicago Plan, sovereign money, and the money-creation literature

The architecture belongs to a well-populated family and should be placed in it honestly. Full-reserve or "Chicago Plan" schemes (Fisher, 1936; Simons, 1936; Benes and Kumhof, 2012) require reserves equal to 100 percent of transaction deposits, so those deposits are no longer created by lending; sovereign-money proposals reserve money creation to the public authority and reduce banks to intermediaries; narrow or limited-purpose banking (Kotlikoff, 2010; Cochrane, 2014) backs demandable claims with safe assets and routes risky lending through equity-funded vehicles. The Citizens Standard sits squarely in the first of these: full-reserve transaction accounts, at-risk term deposits, a payment system insulated from the credit cycle. Its single departure is distributional. The 100-percent-money literature almost uniformly routes the seigniorage from reclaimed money creation to the government — through spending, tax reduction, or debt cancellation — and treats distribution as secondary. The Citizens Standard distributes it equally to citizens as locked individual equity (Neo-Solon, 2026a, §17.3, §17.4). The banking reform is borrowed; the distribution is the novel element.

This placement also fixes what this paper must establish. Because the framework adopts full reserve, it inherits the tradition's claimed benefits — bounded money, no runs on the payment system, a credit cycle severed from the money supply — but it also inherits the tradition's open problems: the transitional credit contraction (Benes and Kumhof, 2012, find it manageable but real) and the boundary problem (Goodhart, 2008), the tendency of highly liquid term claims to migrate into transactional use. The framework does not claim to have dissolved these; it claims a specific, bounded handling of each, set out in Sections 5 and 7.

Relative to the closest prior work, this is a paper of integration rather than re-evaluation. Benes and Kumhof (2012) supply the calibrated macroeconomic case that a full-reserve transition is feasible and can dampen the credit cycle; they do not address what full reserve does inside a two-circuit architecture, nor where the reclaimed seigniorage goes. This paper adds three things to that base: it shows that full reserve hands the two-circuit model its monetary-control and determinacy results outright rather than as a defended bound (Sections 3–4); it specifies a credit-supply mechanism in which the conversion contraction is offset by the framework's own citizen issuance channels rather than by a discretionary facility (Section 5); and it routes the

seigniorage to citizens as locked equity rather than to the state. The full-reserve banking is theirs; the integration, the credit-offset mechanism, and the distribution are this framework's.

3. Complete monetary control

Under full reserve the price-level result of the macroeconomic model holds in its strongest form, and it holds because a complication is absent rather than managed. The price level is set by transactional money, $M_T \cdot V/Y$; since $M_T = M_o$, the inflation-gap throttle KI, which adjusts outside transactional money to track a target price path, controls the price-relevant aggregate directly. There is no inside-money term to read, to offset, or to chase. A fractional-reserve version of the framework would have to make the throttle target the sum of outside and bank-created transactional money and reduce its own issuance one-for-one as banks created deposits — workable, but a standing offset operation with a saturation point where deposit creation would exhaust the entire target. Full reserve removes the operation and the saturation point together: there is nothing to offset and no point at which the authority would have to retire more outside money than exists.

Determinacy carries over directly. The forward-looking determinacy of the macroeconomic model (Neo-Solon, 2026e, Proposition 7) was derived for the transactional aggregate. Because that aggregate is now simply M_o , the result applies to the authority's own instrument with no augmented aggregate and no change of form: the rational-expectations equilibrium price path is unique under the money-quantity anchor, the explosive root governing determinacy is unchanged, and the anchor pins the price level without a Taylor-type interest-rate response. The economic content is that full reserve hands the framework its determinacy result outright. In the incumbent system the quantity of transactional money is jointly determined by the authority and by the lending decisions of thousands of private banks; under the Citizens Standard it is set by the authority alone, which is the condition the determinacy proof assumes. This is Proposition N1 (Appendix A.2; proof and Figure 1 in A.3).

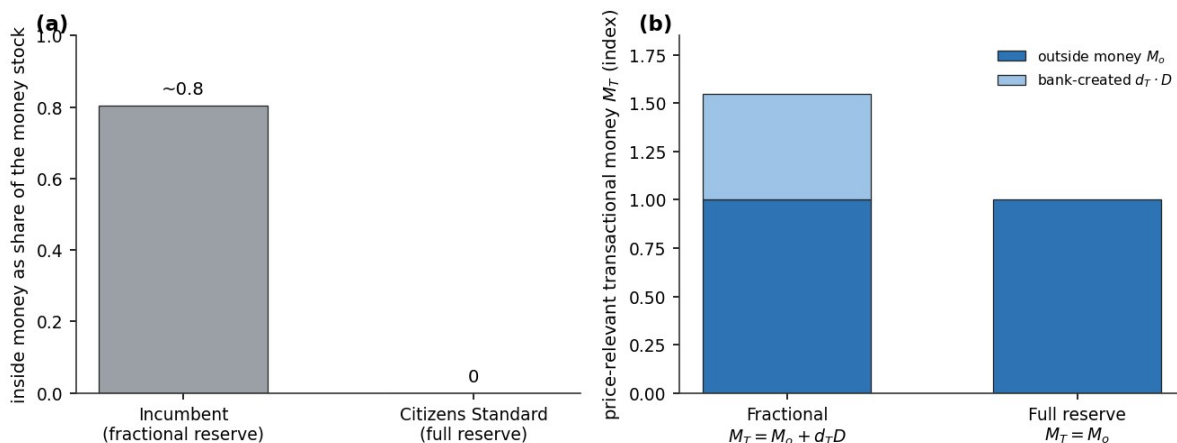


Figure 1. Banks create no transactional money under the Citizens Standard. (a) the inside-money share of the money stock — the share that is a private bank liability — roughly four-fifths in the incumbent system against zero under full reserve; (b) the price-relevant aggregate M_T , which equals the sum of outside and inside transactional money in a fractional system and collapses to outside money M_o alone under full reserve, so the throttle controls it directly.

4. Circuit separation survives bank credit

Removing money creation does not remove bank credit, and credit raises its own version of the separation question. The macroeconomic model already posed and answered it (Neo-Solon, 2026e, §3.8, Proposition 9): if banks lend against asset-circuit wealth and the proceeds are spent on goods, credit opens a channel from the asset circuit into the transactional circuit, re-mixing what the architecture separates. That result is stated there for a system in which credit is money-creating, but its mechanism is about credit-financed *spending*, not about money creation, so it carries over to a full-reserve system unchanged: whether the credit is newly created money or intermediated term deposits, what couples the circuits is the share of credit extended against pledgeable asset wealth and spent into the goods market.

The lock is decisive, here as elsewhere. Let banks lend against pledgeable asset wealth at loan-to-value m , with a fraction ϕ_{liq} of asset wealth liquid and a share χ_c of new credit spent on goods. The induced asset-to-consumer coupling is $\chi_c \cdot m \cdot \phi_{liq}$, and writing the credit intensity as $\kappa_{bank} \equiv m \cdot \phi_{liq}$, circuit separation survives so long as the total coupling $\lambda_{leak} + \chi_c \cdot \kappa_{bank}$ stays below the determinacy threshold $\zeta \approx 0.13$ — *that is, while $\kappa_{bank} < (\zeta - \lambda_{leak})/\chi_c \approx 0.32$* under the conservative, financial-accelerator reading (Neo-Solon, 2026e, A.12). The locked floor settles the question. Because the dominant household asset is non-tradeable and therefore non-pledgeable, only the liquid fraction backs credit, $\phi_{liq} \approx 0.15$, so the baseline credit intensity is $\kappa_{bank} \approx 0.075$ — far inside the threshold. Breaking separation through bank credit would require a loan-to-value above two, which is not a feasible loan. Were the floor fully pledgeable, the same calculation would break separation at an ordinary mortgage loan-to-value of about 0.32; the non-pledgeability of the floor is exactly what prevents it.

Full reserve, if anything, strengthens the conclusion. The coupling above runs entirely through credit-financed spending out of pledged wealth; it does not depend on, and is not amplified by, money creation, because under full reserve there is none. The rule-bound structural buyer supplies an independent second safeguard — because asset demand is dominated by the K1/K2 buyer rather than by credit, the financial accelerator is damped and the critical intensity rises further (Neo-Solon, 2026e, Proposition 9, part iv) — but the lock alone already places the baseline comfortably inside the separation region. This is Proposition N2 (Appendix A.2).

5. The Loanable-Funds Frontier under Full Reserve

The substantive cost of full reserve is not monetary control or separation, both of which it secures; it is credit supply. When banks may lend only what savers term out, the elasticity that money creation provided is gone, and the framework must show both that credit can be supplied at adequate scale and that the transition to full reserve does not impose a credit collapse. This section sets out the steady-state identity, locates what full reserve actually binds, accounts for the conversion contraction, and characterizes the residual steady-state cost.

5.1 The loanable-funds identity and the leverage cap

Under full reserve, bank credit is funded entirely by term-deposit liabilities and equity, and capped by leverage. With term deposits D and equity E and a maximum leverage of four to one, credit assets satisfy $L \leq 4E$ and $D \leq 3E$, so a term-deposit base of D supports bank credit of up to $D + E = D \cdot 4/3$. At the framework's launch calibration the relevant magnitudes are concrete: with a term-deposit share near 60 percent of launch M2 (Neo-Solon, 2026c), the term-deposit

base is about \$13.4 trillion, the implied equity is about \$4.5 trillion, and bank credit capacity is about \$17.9 trillion. The leverage cap is countercyclical: it tightens to three to one when private credit growth exceeds nominal output growth by more than three points for two quarters, and loosens to five to one in the opposite case (Neo-Solon, 2026a, §9.2), so the credit system has a built-in brake on the procyclicality that money-creating banks exhibit.

5.2 Full reserve binds only the bank slice

A common objection to full reserve is that it would starve the economy of credit. The objection mistakes the scope of the constraint. Full reserve binds only credit that would otherwise be funded by money creation — the bank-intermediated slice. Credit supplied directly through markets — corporate bonds, commercial paper, and other non-bank channels — does not involve deposit creation in the first place and is untouched by the reserve requirement. In the United States, private non-financial credit outstanding is about \$42.4 trillion (Federal Reserve Z.1, household \$20.5 trillion plus business \$21.9 trillion), of which the bank-funded share that full reserve actually constrains is a minority; large corporate borrowers in particular finance through markets that never depended on bank deposit funding. The full-reserve constraint therefore bites on a bounded slice of total credit, not on its whole, and the question of Section 5.3 is how that slice is re-funded, not how the entire credit stock is replaced.

5.3 The conversion contraction and its sovereign offset

The transition to full reserve removes bank money-creation capacity, which is contractionary on the bank slice while it occurs, and potentially deflationary. The framework offsets this with its own sovereign issuance, routed through the channels that already exist to distribute new money to citizens — not through a discretionary central-bank facility. Two instruments carry the load during the transition (Neo-Solon, 2026c). The Transition Lending Facility channels a bounded share — at most 30 percent of K2 seigniorage — into zero-spread replacement lending for the most bank-dependent borrowers, covering an estimated 12 to 38 percent of the annual credit at risk; the KT channel, which directs sovereign issuance to debt redemption during the same window, supplies outside money that offsets a further 41 to 59 percent. The two together cover roughly 53 to 97 percent of the credit at risk, and the residual is absorbed by the conditional KI_T damper (Neo-Solon, 2026e, §4.7), a non-binding-by-default channel that issues additive spendable money only when the residual contraction would pull the price path below target, and self-extinguishes once banking separation completes. The point is structural: the same sovereign issuance that removes inside-money creation, in the form of full-reserve conversion, simultaneously supplies the outside money that fills the gap — the contraction and its offset are two sides of one operation. The framework does not claim the offset is complete at every moment; it claims a better failure mode and a materially reduced magnitude, with the residual handled by a rule whose stability is inherited from the convergence result of the macroeconomic model (Neo-Solon, 2026e, Proposition 3).

5.4 Steady-state credit elasticity and the countercyclical channels

After conversion, a residual cost remains, and it is the honest one. Money creation gave the incumbent system an elastic credit supply: banks could expand credit without first attracting savings. Full reserve removes that elasticity, so the steady-state bank credit supply rides the public's willingness to hold term deposits — and the interest-elasticity of saving is, empirically, small and of uncertain sign, with the uncompensated elasticity typically estimated near zero

(Bernheim, 2002). The consequence is that a shortfall in term-deposit funding cannot be cleared cheaply by a higher deposit rate, because saving barely responds; it clears instead by a higher loan rate that rations credit — a crunch — or by sovereign replacement lending. Figure 2 makes the trade explicit on the bank slice: under a 30-percent term-deposit shortfall, holding the loan rate steady requires the sovereign channels to fund about 30 percent of the affected credit, while letting the rate clear the shortfall on price alone would take a loan-rate premium on the order of six points and contract credit by roughly a quarter. The framework's answer is the countercyclical machinery it already has — the conditional KI_T damper and the loosening of the leverage cap to five to one in a credit contraction — which restores the elasticity money creation used to supply, without restoring money creation. The genuine cost of full reserve is thus not a permanent credit shortage but a dependence on the sovereign channels to supply the cyclical elasticity that private money creation supplied before; the framework internalizes that dependence rather than denying it. This is Proposition N3 (Appendix A.2; Figure 2 in A.3).

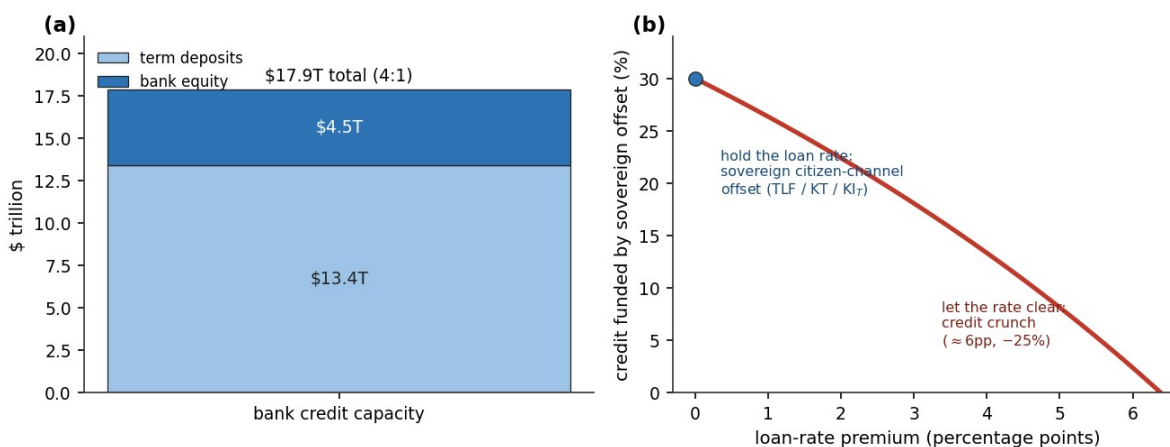


Figure 2. The loanable-funds frontier under full reserve. (a) bank credit capacity as term deposits plus equity under the four-to-one leverage cap, about \$17.9 trillion at the launch split; (b) the credit-clearing frontier under a term-deposit shortfall: the affected credit is funded either by the sovereign citizen-channel offset (holding the loan rate steady) or by a rising loan rate that rations credit (a crunch), with saving's near-zero rate-elasticity making the price route costly. Calibration: term share 0.60 of launch M2, leverage 4:1, saving semi-elasticity 0.01/pp, credit semi-elasticity 0.04/pp. Robustness. Because the frontier uses fixed elasticities, the replication sweeps them from low (both halved) through baseline to high (both doubled) under the same 30 percent shortfall. The rate premium moves with the elasticities (from about 3 to 13 percentage points, larger where markets are less elastic, as expected), but a material credit crunch of roughly one quarter of bank credit appears in every case, and the empirical evidence places saving near the inelastic end where the crunch is largest. The qualitative result — that clearing a term-deposit shortfall on price alone rations credit rather than conjuring saving, so the citizen-channel offset and the countercyclical leverage rule are the design's answer — is therefore not an artifact of one calibration.

6. Payment–credit separation and run-proofness

A banking sector ordinarily introduces fragility absent from a pure outside-money system: deposits are redeemable on demand, and a loss of confidence can convert them to base money

en masse (Diamond and Dybvig, 1983). Full reserve removes that fragility at its source, by construction rather than by insurance, and the locked floor adds a further layer.

Three layers, the first decisive. Transaction accounts are held one-for-one against reserves, so they cannot be run: there is no maturity or reserve mismatch to expose, and a panic that converted every transaction balance to currency would find the reserves already there. Term deposits are at-risk term claims, not demandable money; a holder cannot withdraw a term deposit at will, so the run technology — mass conversion of demandable claims — does not apply to them either. They can lose value if the bank's loans sour, but a loss borne by an explicit risk-taker at maturity is a solvency event, not a liquidity run. And the dominant household asset, the locked floor, is equity outside the banking system altogether: it is neither demandable nor pledgeable, so it cannot be run and does not appear in the runnable base at all.

The bound. The maximum systemic money contraction is therefore bounded by the term-deposit share of the money stock — on the order of the launch term share rather than the whole deposit base. The contrast with 1930–1933, when essentially all deposits were simultaneously runnable and the money supply collapsed, is the point: under the Citizens Standard the fully reserved transaction layer cannot contract at all, so the worst case is confined to losses on the at-risk term layer (Neo-Solon, 2026c, banking-separation analysis). This does not eliminate the case for ordinary supervision or resolution of the term-deposit layer; it bounds the scale of what supervision must cover, and removes the payment system from the failure mode entirely. The architecture's contribution to stability is structural rather than discretionary: it does not insure the run, it removes the run technology from the core of the balance sheet. This is the Fisher (1936) and Benes and Kumhof (2012) no-runs result, here resting on the reserved transaction layer first and the locked floor second. This is Proposition N4 (Appendix A.2; Figure 3 in A.3).

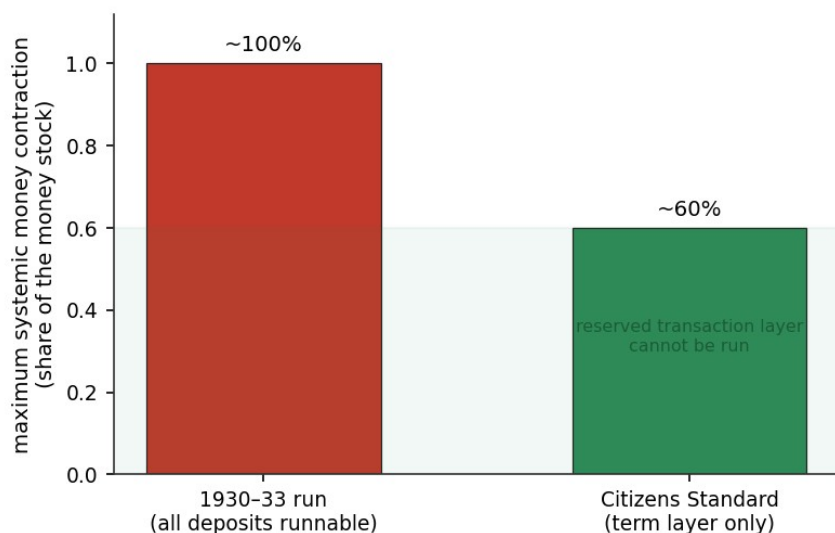


Figure 3. Run-proof payments bound the loss. The maximum systemic money contraction: essentially the entire money stock in a fractional-reserve run such as 1930–1933, against the term-deposit share alone under the Citizens Standard, because the reserved transaction layer cannot be run and the locked floor lies outside the runnable base.

7. The boundary problem: near-money

The framework should concede, rather than minimize, the standing critique of narrow and full-reserve banking: the boundary problem (Goodhart, 2008). Highly liquid term claims tend to migrate into transactional use — a money-market instrument that settles like a deposit becomes, in practice, a near-money — so the separation of money from non-money is never as clean in operation as on the statute book. A full-reserve rule does not abolish this tendency; it raises the cost and visibility of near-money creation without ending it.

The Citizens Standard's handling is partial and should be stated as such. Its one structural advantage is that the throttle targets the *total* transactional aggregate, not the authority's issuance alone: if a class of near-money comes to circulate as means of payment, it enters the price-relevant quantity, and the throttle offsets it exactly as it would any other transactional money — the one place the offset logic of a fractional system survives, now applied to near-money rather than to bank deposits. The binding qualification is observability. The offset works only on near-money the authority can see and measure; near-money that circulates below the measurement threshold escapes it. The honest claim is therefore conditional rather than absolute: full reserve plus a total-aggregate throttle contains observable near-money, and the framework's falsifiable commitment is to monitor the boundary — to track the migration of liquid term claims into transactional use and to bring observed near-money into the targeted aggregate — not to assert that the boundary holds by construction. This is Proposition N5 (Appendix A.2).

8. Discussion and relation to the series

What the reduced form omits. The banking sector here is summarized by a small set of sufficient statistics — the term-deposit share, the loan-to-value ratio, the pledgeable fraction, the leverage cap, and the sovereign-offset coverage — rather than by a fully articulated intermediary with an optimizing balance sheet. This is deliberate: the questions the paper asks turn on aggregate quantities these statistics capture. A heavier apparatus is required for questions the reduced form cannot reach — the dynamics of credit over the cycle, the pricing of bank risk, and the welfare comparison between bank intermediation and the framework's own issuance channels. A microfounded balance-sheet model in the tradition of Gertler and Kiyotaki (2010) is the natural successor and is left to future work. The normative question — whether, and how much, bank intermediation improves capital allocation relative to direct market credit and the K-channels — is genuinely separate and is not settled here; a thin welfare claim would be the weakest part of the paper, and the deferral is honest rather than evasive.

Relation to the rest of the series. The results cohere with the framework's other papers through two devices it already owns. The lock — the non-pledgeability of the floor — is why bank credit cannot re-mix the circuits (Section 4) and is the third layer of run-proofness (Section 6); it is the same device behind the demographic, labor, and separation results elsewhere. The sovereign issuance engine — the K-channels and the conditional KI_T damper — is what offsets the conversion contraction and supplies the cyclical credit elasticity full reserve removes (Section 5); it is the same engine that distributes new money to citizens. The architectural paper specifies the full-reserve banking design and its lineage (Neo-Solon, 2026a, §9, §17.3); the transition paper sequences the banking separation and sizes its offsets (Neo-Solon, 2026c, Phase 3); the macroeconomic model supplies the determinacy and separation results this paper

inherits (Neo-Solon, 2026e, Propositions 7 and 9). This paper's role is to show that those pieces compose into a coherent full-reserve banking layer — that the framework's banking claims are consequences of structure it already has.

9. Conclusion

The Citizens Standard is a full-reserve system, and its banking layer should be read as a member of the Chicago-Plan tradition with a distributional departure, not as a defense of bounded private money creation. Within that reading the results are clean. Monetary control is complete because banks create no transactional money, so the price-relevant aggregate is the authority's own instrument and the determinacy result applies outright. Circuit separation survives bank credit because the only channel credit could open is bounded by the non-pledgeability of the locked floor, far inside the threshold. Credit is supplied by genuine intermediation of term deposits and equity under a countercyclical leverage cap; full reserve binds only the bank slice, the conversion contraction is offset by the framework's own sovereign channels to near-full coverage, and the residual steady-state cost — the loss of credit elasticity — is restored by those same channels without restoring money creation. The payment system is run-proof because it is fully reserved, with at-risk term deposits and locked equity as further layers, so the worst-case money contraction is bounded by the term-deposit share rather than the whole. And the boundary problem is conceded honestly: full reserve raises the cost and visibility of near-money without abolishing it, with the total-aggregate throttle containing what it can observe. The banking reform is borrowed from a tradition with a serious pedigree; what the framework adds is where the money goes.

Appendix

A.1 Notation

Symbol	Meaning
M_o	Outside transactional money (authority-issued, KI-controlled)
M_A	Asset-circuit money
$M_T = M_o$	Effective transactional money under full reserve; $P = M_T \cdot V/Y$
D	Term deposits — at-risk investment claims, not money
E	Bank equity
L	Bank loans, with $L \leq 4E$ and $D \leq 3E$ under the leverage cap
m	Loan-to-value ratio
ϕ_{liq}	Liquid (pledgeable) fraction of asset wealth; $(1 - \phi_{liq})$ locked
$\kappa_{bank} \equiv m \cdot \phi_{liq}$	Credit intensity of bank lending
χ_c	Share of new credit spent into the

Symbol	Meaning
	transactional circuit
$\lambda_{\text{leak}} \approx 0.03$	Structural asset-to-consumer leak
$\zeta \approx 0.13$	Separation / determinacy threshold (Neo-Solon, 2026e, A.10, A.12)
TLF	Transition Lending Facility (≤ 0.30 of K2 seigniorage)
KT	Transition redemption channel
KI_T	Conditional transition damper

A.2 Propositions

Proposition N1 (complete monetary control). Under full reserve, banks create no transactional money, so the price-relevant aggregate is $M_T = M_o$. Then (i) the throttle KI controls the price-relevant quantity directly, with no inside-money term to offset and no saturation point; and (ii) the forward-looking determinacy of Proposition 7 (Neo-Solon, 2026e) applies to M_o without an augmented aggregate, so the equilibrium price path is unique under the money-quantity anchor. Proof and Figure 1 in A.3.

Proposition N2 (separation survives bank credit). Let banks lend against pledgeable asset wealth at loan-to-value m , a fraction ϕ_{liq} of asset wealth liquid, a share χ_c of credit spent on goods, with credit intensity $\kappa_{\text{bank}} \equiv m \cdot \phi_{\text{liq}}$. Then the asset-to-consumer coupling is $\lambda_{\text{leak}} + \chi_c \cdot \kappa_{\text{bank}}$, independent of whether credit is money-creating or intermediated; circuit separation holds iff this stays below $\zeta \approx 0.13$, i.e. $\kappa_{\text{bank}} < (\zeta - \lambda_{\text{leak}}) / \chi_c \approx 0.32$ (conservative reading); and because the floor is non-pledgeable, $\phi_{\text{liq}} \approx 0.15$ and $\kappa_{\text{bank}} \approx 0.075$ lie far inside, with the critical loan-to-value $m = \kappa_{\text{bank}} / \phi_{\text{liq}} > 2$ infeasible. Proof in A.3; the underlying result is Proposition 9 of Neo-Solon (2026e).

Proposition N3 (the Loanable-Funds Frontier under Full Reserve). Bank credit satisfies $L \leq D + E$ with $D \leq 3E$ (leverage $\leq 4:1$), so a term-deposit base D supports credit up to $D \cdot 4/3$. (i) Full reserve binds only bank-intermediated credit; market credit is unaffected. (ii) The conversion contraction on the bank slice is offset by sovereign issuance — TLF (≤ 0.30 of K2 seigniorage; 12–38% coverage) plus KT (41–59%) plus the residual KI_T damper — approaching full coverage. (iii) In steady state, because the interest-elasticity of saving is ≈ 0 , a term-deposit shortfall clears either by a rationing rise in the loan rate or by sovereign replacement lending; the countercyclical KI_T damper and the loosening of leverage to 5:1 restore the lost elasticity without restoring money creation. Figure 2 in A.3.

Proposition N4 (run-proof payments). Fully reserved transaction accounts cannot be run (no maturity or reserve mismatch); at-risk term deposits are not demandable money, so the Diamond–Dybvig run technology does not apply to them; the locked floor is non-demandable equity outside the runnable base. Hence the maximum systemic money contraction is bounded by the term-deposit share of the money stock, not the whole deposit base, and the lender-of-last-resort need is confined to the at-risk term layer. Proof and Figure 3 in A.3.

Proposition N5 (the boundary problem). A full-reserve rule raises the cost and visibility of near-money creation without abolishing it (Goodhart, 2008). Because the throttle targets the total transactional aggregate, observable near-money that circulates as means of payment enters the price-relevant quantity and is offset one-for-one. The containment is therefore conditional on observability; the framework's commitment is to monitor the boundary and bring observed near-money into the targeted aggregate, not to assert the boundary holds by construction.

A.3 Proofs

N1. The price equation $P = M_T \cdot V/Y$ with $M_T = M_o$ (no bank-created transactional money under full reserve) gives $P = M_o \cdot V/Y$, so KI, setting M_o , sets P directly; there is no D term, hence no offset $M_o = M - d_T \cdot D$ and no saturation condition $d_T \cdot D \leq M$. For (ii), the equilibrium conditions of Proposition 7 (Neo-Solon, 2026e) depend on transactional money; with that aggregate equal to M_o , the Cagan-type demand and the quantity rule are unchanged in form, the explosive root and the determinacy conclusion are unchanged, and no augmented aggregate is introduced. ■

N2. Banks lend against pledgeable wealth, capacity $m \cdot \phi_{liq} \cdot A$ with A asset-circuit wealth and the locked floor excluded by non-pledgeability; a share χ_c of new credit is spent on goods, inducing a transactional inflow $\chi_c \cdot m \cdot \phi_{liq}$ per unit of asset wealth. Adding the structural leak gives total coupling $\zeta_{eff} = \lambda_{leak} + \chi_c \cdot \kappa_{bank}$ with $\kappa_{bank} \equiv m \cdot \phi_{liq}$. The derivation is identical whether the credit is created money or intermediated term deposits, since only the spent share enters; full reserve removes a money-creation amplification that would otherwise raise, not lower, the coupling. Inserting ζ_{eff} into the determinacy system of Proposition 7 fails determinacy at $\zeta \approx 0.13$ (A.10), so separation requires $\kappa_{bank} < (\zeta - \lambda_{leak})/\chi_c \approx 0.32$ at $\chi_c = 0.30$, $\lambda_{leak} = 0.03$. With $\phi_{liq} \approx 0.15$, $\kappa_{bank} \approx 0.075$ and $m = \kappa_{bank}/\phi_{liq} \approx 2.1$, infeasible; fully pledgeable wealth ($\phi_{liq} = 1$) gives $m^* \approx 0.32$. ■

N3. Leverage $\leq 4:1$ gives $L \leq 4E$ and $D \leq 3E$, so credit $\leq D + E = D \cdot 4/3$; at $D = 0.60 \cdot M2 \approx \$13.4T$, $E \approx \$4.5T$ and $L \leq \$17.9T$. (i) Only deposit-funded credit is constrained by the reserve rule; market credit creates no deposits and is unaffected. (ii) Conversion-contraction coverage is TLF ($\leq 0.30 \cdot K2$ seigniorage, 12–38%) + KT (41–59%) = 53–97%, residual by KI_T (Neo-Solon, 2026c, 2026e §4.7). (iii) With saving supply $S(r) = S_o(1 + \varepsilon_s \cdot (r - r_o))$, $\varepsilon_s \approx 0$, and credit demand $C(r) = C_o(1 - \varepsilon_d \cdot (r - r_o))$, $\varepsilon_d \approx 0.04$, a shortfall $S_o < C_o$ clears at $r = r_o + (C_o - S_o)/(S_o \cdot \varepsilon_s + C_o \cdot \varepsilon_d)$; at a 30% shortfall on the bank slice, $r - r_o \approx 6$ points with credit contracting $\approx 25\%$, versus a sovereign offset of $\approx 30\%$ holding the rate. Because $\varepsilon_s \approx 0$, the price route is predominantly rationing. ■

N4. A run converts demandable claims to base money. Reserved transaction accounts are backed one-for-one, so conversion finds the reserves in place and the runnable base excludes them. Term deposits are non-demandable within term, so they are excluded from the run technology (a maturity loss is a solvency event, not a run). The locked floor is non-demandable and non-pledgeable, excluded by construction. The runnable base is therefore at most the term-deposit share of the money stock, and the maximum contraction and lender-of-last-resort need scale with it. ■

N5. Statement is a containment-plus-condition, not an identity. Near-money that circulates as means of payment enters M_T by definition of the transactional aggregate; the throttle, targeting

M_T, offsets it one-for-one, which is the N1 mechanism applied to near-money. The offset operates only on the measured aggregate, so unobserved near-money is not contained; the result is the conditional statement, with observability the falsifiable margin (Goodhart, 2008). ■

A.4 Calibration and falsification

Term-deposit share and the loanable base. The share of M2 citizens hold as at-risk term deposits rather than reserved transaction accounts sets the bank credit base (N3) and the run-proofness bound (N4); the launch value (≈ 0.60) is a framework split, estimable from account data after enactment, and the credit-supply and run results are monotone in it.

Pledgeable fraction (φ_{liq}). The liquid, pledgeable share of asset wealth — held well below one by the locked floor — sets the credit intensity κ_{bank} and hence the separation margin (N2); estimating φ_{liq} below the threshold that keeps κ_{bank} under ζ^* is the falsifiable content, inherited from Neo-Solon (2026e, A.12).

Sovereign-offset coverage (TLF, KT, KI_T). Whether the citizen-channel offset covers the conversion contraction is the transition's central empirical question; the framework's prediction is 53–97% from TLF and KT with the residual absorbed by KI_T, falsifiable against realized credit-at-risk and issuance during a phased separation (Neo-Solon, 2026c).

The boundary (near-money). The migration of liquid term claims into transactional use (N5) is the standing falsification risk; the commitment is to measure it and fold observed near-money into the targeted aggregate. A regime in which near-money grows faster than it can be observed and offset would defeat the containment — the honest failure mode of the design.

References

- Benes, J., and Kumhof, M. (2012). The Chicago Plan revisited. IMF Working Paper 12/202.
- Bernheim, B. D. (2002). Taxation and saving. In A. J. Auerbach and M. Feldstein (eds.), *Handbook of Public Economics*, vol. 3, ch. 18, 1173–1249. Amsterdam: Elsevier.
- Cochrane, J. H. (2014). Toward a run-free financial system. In M. N. Baily and J. B. Taylor (eds.), *Across the Great Divide: New Perspectives on the Financial Crisis*. Hoover Institution Press.
- Diamond, D. W., and Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3), 401–419.
- Fisher, I. (1936). *100% Money*. New York: Adelphi.
- Friedman, M. (1948). A monetary and fiscal framework for economic stability. *American Economic Review*, 38(3), 245–264.
- Gertler, M., and Kiyotaki, N. (2010). Financial intermediation and credit policy in business cycle analysis. In *Handbook of Monetary Economics*, vol. 3, 547–599.
- Goodhart, C. A. E. (2008). The boundary problem in financial regulation. *National Institute Economic Review*, 206, 48–55.
- Jakab, Z., and Kumhof, M. (2015). Banks are not intermediaries of loanable funds — and why this matters. Bank of England Working Paper 529.
- Kotlikoff, L. J. (2010). *Jimmy Stewart Is Dead: Ending the World's Ongoing Financial Plague with Limited Purpose Banking*. Hoboken, NJ: Wiley.
- McLeay, M., Radia, A., and Thomas, R. (2014). Money creation in the modern economy. Bank of England Quarterly Bulletin, 2014 Q1.
- Simons, H. C. (1936). Rules versus authorities in monetary policy. *Journal of Political Economy*, 44(1), 1–30.
- Tobin, J. (1963). Commercial banks as creators of money. In *Banking and Monetary Studies*, ed. D. Carson. Homewood, IL: Richard D. Irwin.
- Tobin, J. (1985). Financial innovation and deregulation in perspective. *Bank of Japan Monetary and Economic Studies*, 3(2), 19–29.
- Neo-Solon (2026a). The Citizens Standard: One Model, Many Systems — A Constitutional Monetary Architecture. SSRN 6702518.
- Neo-Solon (2026b). The Citizens Standard: A Historical Counterfactual. SSRN 6735078.
- Neo-Solon (2026c). The Citizens Standard: Transition Architecture and Migration Mechanics. SSRN 6810741.
- Neo-Solon (2026d). The Citizens Standard: A Statutory Implementation Pathway. SSRN 6873798.
- Neo-Solon (2026e). The Citizens Standard: A Macroeconomic Model of a Two-Circuit Monetary System.