Monetary Production Theory and the Threshold Structure of Financial Crises

Francisco Borja Avilés*

February 27, 2026

Abstract

I propose a theory of financial crises based on the concept of *monetary production*: the technology that transforms backing assets into liquid media accepted as money. Reinterpreting the general liquidity constraint of Brunnermeier and Niepelt (2019) as a production technology, I derive a computable crisis threshold $\tilde{q}^* = (1 - \beta - \omega)/(\beta\kappa)$, determined by the discount factor β, the effective liquidity intensity $\omega = \lambda\bar{v}$ (combining the nominal interest rate and the velocity weight of the liquid media), and the backing-asset cash-flow share κ. Below $\tilde{q}^*$, the self-referential loop sustaining inside money—liquid media are accepted because the issuer is solvent, and the issuer is solvent partly because its liquid media are accepted—collapses through a saddle-node bifurcation. The threshold partitions the parameter space into three regimes based on exogenous fundamental asset quality: a safe region in which the monetary equilibrium is unique; a fragile region in which both monetary and non-monetary equilibria coexist, so that sunspot-driven crises are possible; and a crisis region in which the only equilibrium features zero issuance of liquid media. The production frontier is increasing and linear, but the equilibrium response is discontinuous at $\tilde{q}^*$: crises are abrupt because the self-sustaining monetary equilibrium ceases to exist, not because of frontier curvature. I demonstrate that banking, currency, sovereign debt, and debasement crises are all instances of this single mechanism. Finally, a quantitative calibration contrasts the 2023 SVB episode with the 1929 U.S. banking sector, revealing two distinct structural regimes: idiosyncratic, avoidable fragility, and system-wide, permanent fragility.

Keywords: financial crises, money creation, liquidity, banking, currency crises, monetary production

JEL Classification: E41, E44, E51, G01, G21

*Universidad San Francisco de Quito. Email: fborja@usfq.edu.ec. ORCID: 0000-0002-0523-6181.

1 Introduction

Between 1930 and 1933, more than nine thousand American banks failed (Friedman and Schwartz, 1963). The proximate cause was not speculative exposure to complex instruments, but substantial deterioration in one of the most conventional assets on bank balance sheets: agricultural and commercial loans (Calomiris and Mason, 2003). As crop prices declined and borrowers defaulted, the realized cash flows on loan portfolios fell below their contractual promises. Once the effective repayment fraction crossed a critical threshold, deposits backed by those loans—claims that had circulated at par in everyday transactions—ceased to function as reliable media of exchange. The money supply contracted sharply, and real output declined on a scale without modern precedent (Friedman and Schwartz, 1963).

The 1930s banking collapse is one of the starkest instances of a phenomenon that has recurred for eight centuries. From the collapse of the Bardi and Peruzzi banking houses in 1340s Florence, through the Mississippi bubble of 1720, the panics of 1873 and 1907, the East Asian crisis of 1997, to the global financial crisis of 2008—financial crises exhibit a remarkably stable pattern. A monetary producer supplies liquid media backed by underlying assets or resources. For a time, these instruments circulate as money: they are accepted at par, used in daily transactions, and treated as safe. Then the fundamental quality of the backing assets deteriorates. The liquid media lose their money-ness. The money supply contracts. The real economy suffers.

Reinhart and Rogoff (2009) document this pattern across 66 countries over 800 years, famously arguing that each generation believes "this time is different"—that new financial innovations have eliminated old risks. Yet their contribution establishes an empirical regularity rather than a structural mechanism. Why do financial crises recur with such persistence across centuries and institutional environments? Why does the exact same sequence—liquid media, backing assets, collapse—reappear despite substantial changes in financial and legal arrangements? Classic historical accounts often emphasize the role of psychological manias and self-fulfilling panics (Kindleberger and Aliber, 2005). While coordination failures are undeniably important, they are not strictly necessary for a crisis to occur. Indeed, episodes driven entirely by observable fundamental deterioration—such as the rational abandonment of deliberately debased commodity coinage—point to a deeper, mechanical limit.

While the literature lacks a unified, formal definition of a financial crisis, this paper proposes a baseline that operates even under perfect rationality: financial crises are failures of *monetary production*—episodes in which the technology that transforms backing assets into liquid media breaks down.

This definition draws a strict structural boundary between a systemic financial crisis and

a pure asset-price deflation. A market crash only threatens the monetary equilibrium if the collapsing assets are load-bearing inputs for the production of circulating transaction media. This structural boundary explains why severe wealth destructions—from the 1637 Dutch Tulip Mania (Goldgar, 2008; Thompson, 2007) to the 2000 Dot-Com crash (Cochrane, 2009) —did not trigger a financial crisis.

I formalize this mechanism by reinterpreting the general liquidity constraint of Brunnermeier and Niepelt (2019) as a production technology—a mapping from backing-asset portfolios to flows of liquidity services. The central result is that this technology has a computable crisis threshold:

$$\tilde{q}^* = \frac{1-\beta-\omega}{\beta\kappa} \tag{1}$$

below which monetary production collapses through a saddle-node bifurcation. The threshold exists because monetary production is *self-referential*: liquid media are accepted as money because their issuer is solvent, and the issuer is solvent partly because its liquid media are accepted—generating a liquidity premium on its backing assets. When exogenous backing-asset quality $\tilde{q}_t$ falls below $\tilde{q}^*$, this loop cannot close. A second threshold,

$$\tilde{q}^{**} = \frac{1-\beta}{\beta\kappa} \tag{2}$$

marks the boundary above which the producer is solvent on fundamentals alone, without any liquidity premium. Together, the two thresholds partition the state space into three regions: a *safe region* ($\tilde{q}_t \geq \tilde{q}^{**}$) with a unique monetary equilibrium; a *fragile region* ($\tilde{q}^* < \tilde{q}_t < \tilde{q}^{**}$) where both monetary and non-monetary equilibria coexist, so that sunspot-driven crises are possible; and a *crisis region* ($\tilde{q}_t \leq \tilde{q}^*$) where the only equilibrium features zero issuance of liquid media. The width of the fragile region, $\mathcal{G} = \tilde{q}^{**} - \tilde{q}^* = \omega/(\beta\kappa)$, is a computable measure of systemic fragility.

The thresholds are computable from three groups of primitives: the discount factor β, the effective liquidity intensity $\omega = \lambda\bar{v}$ (combining the nominal interest rate i through $\lambda = i/(1+i)$ and the velocity weight $\bar{v}$ of the liquid media), and the cash-flow share κ of backing assets relative to output. They are invariant to the institutional form of money: a bank issuing deposits against mortgages and a currency board holding foreign reserves face the same threshold structure with different input parameters.

This result has two immediate implications. First, banking crises, currency crises, sovereign debt crises, and debasement crises are all instances of the single threshold condition $\tilde{q}_t < \tilde{q}^*$, applied to monetary arrangements with different backing resources—unifying the cross-century regularity of Reinhart and Rogoff (2009) within one mechanism. Second, the "this time is

different" syndrome is the belief that $\tilde{q}^*$ has shifted when it has not: each era's monetary innovators believe their new form of backing is immune to quality deterioration, but $\tilde{q}^*$ depends on primitives that do not vary with institutional innovation.

The quantitative calibration illustrates these two regimes. Applied to the 1929 U.S. banking sector—where subsequent failures were driven by observable loan defaults, the precise real-world counterpart of $\tilde{q}_t$—the model yields $\tilde{q}^{**} \approx 1.389 > 1$: the safe region was structurally unreachable, and the system resided in a regime of permanent fragility sustained entirely by the endogenous liquidity premium. The predicted insolvency threshold $\tilde{q}^* \approx 0.854$ aligns with historical recovery rates of 75–80 cents on the dollar for liquidated banks. By contrast, mapping the 2023 Silicon Valley Bank episode into the framework as a hypothetical cash-flow impairment of comparable magnitude places SVB in the fragile region ($\tilde{q}^* < \tilde{q} < \tilde{q}^{**}$ with $\tilde{q}^{**} < 1$)—an idiosyncratic, avoidable vulnerability rather than a systemic structural trap. The two cases demonstrate distinct failure modes within the same three-region state space.

The paper relates to four strands in the literature. First, the monetary-economics literature on liquidity and asset pricing: Brunnermeier and Niepelt (2019) provide the general asset-pricing framework for liquid securities; Lagos and Wright (2005) and Williamson (2012) develop search-theoretic foundations; Nagel (2016) measures the convenience yield on safe assets. While BN utilize their framework to establish conditions for monetary equivalence *conditional on the equilibrium existing*, this paper reinterprets their core liquidity constraint as a production technology to derive the *failure conditions*—the exact threshold below which the monetary equilibrium ceases to exist—and the three-region equilibrium structure.

Second, the crisis literature: Diamond and Dybvig (1983) model bank runs as coordination failures; Bernanke et al. (1999) develop the financial accelerator; Caballero et al. (2017) emphasize safe-asset scarcity; Gorton and Metrick (2012) and Gorton and Ordoñez (2014) analyze collateral-driven crises in the shadow banking system; Dang et al. (2012) develop the information-insensitive debt framework; Reinhart and Rogoff (2009) provide the empirical panorama.

The monetary production framework generates predictions distinct from each of these. Relative to Diamond-Dybvig, crises can occur *without* coordination failure when $\tilde{q}_t$ crosses $\tilde{q}^*$ through gradual asset deterioration—a structural collapse driven by fundamentals, not a sunspot. Relative to the financial accelerator, the causal chain runs from monetary contraction to output decline (monetary channel), not from credit-spread widening to investment decline (credit channel). Relative to safe-asset scarcity, the crisis manifests as a contraction in the *quantity* of liquid media produced, not as an increase in the *price* of existing safe assets. Relative to the information-sensitivity framework, crises in this model do not require

asymmetric information or adverse selection; the monetary equilibrium collapses even when the fundamental deterioration of backing assets ($\tilde{q}_t$) is perfectly observable, because the objective technology of liquidity production can no longer sustain the required balance-sheet solvency.[1]

Third, the macro-finance literature on financial amplification: Brunnermeier and Sannikov (2016) and He and Krishnamurthy (2013) develop models of intermediary-driven amplification through asset-price spirals. The MPT amplification mechanism operates through a distinct *monetary supply* channel: the collapse of inside money reduces the transaction capacity of the economy, which raises the shadow cost of liquidity and further depresses asset prices. The two channels are complementary and can be nested within the general BN framework.

Fourth, the literature on the institutional architecture of money creation: McLeay et al. (2014) describe the operational reality of modern inside money creation, while the free-banking literature (White, 1984; Selgin, 1988) analyzes decentralized monetary production and competitive note issue. This paper formalizes these disparate institutional observations into a single mathematical technology, demonstrating that the limits to money creation—whether by a modern commercial bank, a classical gold-standard note issuance, or a fiat sovereign—are governed by the exact same mechanism.

Section 2 presents the general Brunnermeier-Niepelt framework and the tractable CIA specification used for closed-form results. Section 3 defines the monetary production technology, derives the production frontier and crisis threshold $\tilde{q}^*$, characterizes the amplification mechanism, and provides a quantitative calibration. Section 4 unifies banking, currency, sovereign debt, and debasement crises as instances of the single threshold condition. Section 5 draws policy implications for prudential regulation, lender-of-last-resort design, and monetary architecture. Section 6 concludes.

2 Framework

I adopt the general monetary economy of Brunnermeier and Niepelt (2019). The framework is designed to nest the widest possible range of monetary arrangements—from metallic standards to free banking, from central banking to dollarization, from traditional deposits to stablecoins—within a single set of constraints. This generality is essential for my purpose: a theory of monetary production must apply to every institutional setting in which liquid media are issued against backing assets.

[1]These distinctions are stated as theoretical implications of the model's causal structure, not as formal comparative propositions. Empirical testing of the implied ordering and magnitude predictions requires identification strategies beyond the scope of this paper.

I present the general framework in this section, then specialize to a tractable case (cash-in-advance with log utility) that admits closed-form results. The main theorems in Section 3 are proved in the tractable case; I state their generality as remarks and extensions.

2.1 Environment

Time is discrete, $t \geq 0$. The economy is populated by unit measures of households $h \in \mathcal{H}$, and banks $b \in \mathcal{B}$, plus a monetary authority. Following Brunnermeier and Niepelt (2019), who establish that the relevant medium-of-exchange restrictions and asset pricing conditions share the identical structure for both households and firms, I abstract from a distinct corporate sector and consider an endowment economy. This consolidation streamlines the exposition of the monetary production technology without any loss of generality.

Commodities are indexed by n, with prices ρ_t^n. Securities are indexed by j, with prices p_t^j and payoffs z_t^j in numeraire units. I use a dot $(\cdot)$ in place of an index to denote a vector across all items in that set. Let $a_t^{j,h}$ denote the holdings of security j by household h at the end of period t, such that $a_t^{\cdot,h}$ represents the household's full portfolio vector.

A distinguished security M serves as *outside money*—the economy's base monetary instrument. Its supply $\bar{M}_t$ is determined by the monetary authority. The institutional identity of M varies across regimes: gold under a metallic standard, central bank reserves under fiat, foreign currency under dollarization. My framework is agnostic about this identity; what matters is that M enters the liquidity constraints of agents.

2.2 Households

Each household h maximizes lifetime utility $U^h(x_{\cdot,h})$, where $x_{\cdot,h}$ is the household's consumption vector and U^h is smooth, concave, and strictly increasing. The maximization is subject to two constraints.

Budget Constraint.

$$\sum_j a_t^{j,h} p_t^j = \sum_j a_{t-1}^{j,h}(p_t^j + z_t^j) - \sum_n x_t^{n,h}\rho_t^n - \tau_t^h \qquad \forall\, t \tag{3}$$

The left side is the value of the new portfolio. The right side consists of payoffs from inherited positions, less commodity purchases, less taxes net of transfers.

Liquidity Constraint.

$$L_t^h\Big(\{a_t^{j,h}p_t^j\}_j,\ \{a_{t-1}^{j,h}(p_t^j+z_t^j)\}_j,\ p,\ x_{\cdot,h},\ \rho\Big)\geq 0 \qquad \forall\, t \tag{4}$$

The liquidity constraint L_t^h is the central modeling device. Its arguments include current and inherited portfolio positions (in market-value terms), security prices, consumption plans, and commodity prices. By specifying L_t^h appropriately, one nests virtually all standard models of money demand: cash-in-advance (Lucas, 1984; Svensson, 1985), money-in-utility-function (Sidrauski, 1967), shopping time (McCallum, 1983), and borrowing constraints (Kiyotaki and Moore, 2012; Holmström and Tirole, 1998).

The key feature for my purposes is that *multiple* securities may enter L_t^h—outside money M, bank deposits D^b, banknotes N^b, stablecoins—each with a potentially different weight reflecting its liquidity value. This is what makes monetary production possible: for example, banks issue liabilities D^b that enter L_t^h with positive weight, backed by assets that may or may not enter L_t^h directly.

2.3 Banks

Banks are owned by households. Bank b maximizes the market value of its dividend stream:

$$\max \sum_t \mathbb{E}_0\Big[\mu_{0,t}\, z_t^b\Big] \tag{5}$$

where $\mu_{t,s}$ is the date-t stochastic discount factor for payoffs at date s, and z_t^b is bank b's dividend.

Bank Budget Constraint.

$$\sum_{j\neq b} a_t^{j,b}p_t^j = \sum_{j\neq b} a_{t-1}^{j,b}(p_t^j+z_t^j) - z_t^b \qquad \forall\, t \tag{6}$$

where $j\neq b$ excludes the bank's own equity from the summation.

Bank Liquidity Constraint.

$$L_t^b\Big(\{a_t^{j,b}p_t^j\}_j,\ p\Big)\geq 0 \qquad \forall\, t \tag{7}$$

This captures reserve requirements, regulatory constraints, and balance sheet restrictions arising from incentive problems (Diamond and Rajan, 2001; Calomiris and Kahn, 1991).

Under different institutional regimes, L_t^b varies: regulated banks face statutory reserve ratios; free banks face market-discipline constraints; shadow banks may face no formal constraint at all (which, as I shall show, does not free them from the production-technology constraint).

2.4 Monetary Authority

The monetary authority controls the supply of outside money $\bar{M}_t$ and levies lump-sum taxes net of transfers $\{\tau_t^h\}_h$ (with $\tau_t^h > 0$ denoting a net tax on household h). Its flow budget constraint is:

$$\sum_{j \notin \mathcal{J}_L^c} a_t^{j,c} p_t^j = \sum_{j \notin \mathcal{J}_L^c} a_{t-1}^{j,c}(p_t^j + z_t^j) + \underbrace{(\bar{M}_t - \bar{M}_{t-1})\, p_t^M}_{\text{seigniorage}} + \int_{\mathcal{H}} \tau_t^h \, dh \qquad \forall\, t \tag{8}$$

where $\mathcal{J}_L^c$ denotes the set of monetary liabilities issued by the authority (i.e., outside money M), excluded from the asset-side summation, and $(\bar{M}_t - \bar{M}_{t-1})p_t^M$ is the seigniorage revenue from net new money issuance. The constraint holds subject to a no-Ponzi game condition. **Crucially for my framework, the authority maintains a balance sheet where the outstanding supply of outside money $\bar{M}_t$ is backed by a portfolio of assets $\{a_t^{j,c}\}$ (which may include foreign exchange reserves, government bonds, or capitalized fiscal transfers). The fundamental value of this portfolio will dictate the quality of outside money, directly paralleling the backing requirements of private banks.**

The authority may be a central bank, a currency board, a government treasury, or absent entirely (as under free banking or full dollarization). What matters for my theory is not the authority's institutional form, **but that it operates as a monetary producer subject to the same technological backing constraints as private institutions.** Whether this authority subsequently expands $\bar{M}_t$ to act as a lender of last resort during banking-sector stress is a policy question I take up in Section 5.

2.5 Equilibrium

Definition 1 (Competitive Equilibrium)**.** Conditional on monetary authority policy $\{\bar{M}_t, \{\tau_t^h\}_h\}$, a competitive equilibrium consists of:

(i) A consumption allocation $\{x^{\cdot,h}\}_h$;

(ii) Price and payoff vectors $(p, q, \{z^j\}_j)$;

(iii) Stochastic discount factor processes $\{\mu_{0,t}^h\}_{t,h}$;

(iv) Portfolio sequences $(\{a^{\cdot,h}\}_h, \{a^{\cdot,b}\}_b)$;

such that all agents solve their respective programs and markets clear:

$$\int_{\mathcal{H}} a_t^{j,h}\, dh + \int_{\mathcal{B}} a_t^{j,b}\, db + a_t^{j,c} = \bar{S}^j \qquad \forall\, j,\, t \tag{9}$$

where $\bar{S}^j$ is the exogenous aggregate supply of security j ($\bar{S}^j = 0$ for financial claims in zero net supply; $\bar{S}^j > 0$ for the real endowment asset backing monetary production).

2.6 Asset Pricing with Liquidity

The household's first-order condition for security j yields the Euler equation:

$$p_t^j = \mathbb{E}_t\Big[\mu_{t,t+1}^h \cdot (p_{t+1}^j + z_{t+1}^j)\Big] + \ell_t^{j,h} \tag{10}$$

where $\mu_{t,t+1}^h$ is the stochastic discount factor and $\ell_t^{j,h}$ is the *liquidity payoff*—the value of relaxing current and future liquidity constraints by holding one additional unit of security j.

Equivalently, defining the *liquidity kernel*:

$$\Lambda_{t,t+1}^{j,h} = \frac{1 + \lambda_{t+1}^h \cdot \frac{\partial L_{t+1}^h}{\partial (a_t^{j,h}(p_{t+1}^j + z_{t+1}^j))}}{1 - \lambda_t^h \cdot \frac{\partial L_t^h}{\partial (a_t^{j,h} p_t^j)}} \tag{11}$$

where $\lambda_t^h = \tilde{\lambda}_t^h / \tilde{\mu}_t^h$ is the normalized shadow price of the liquidity constraint, the Euler equation becomes:

$$p_t^j = \mathbb{E}_t\Big[\mu_{t,t+1}^h \cdot \Lambda_{t,t+1}^{j,h} \cdot (p_{t+1}^j + z_{t+1}^j)\Big] \tag{12}$$

The liquidity kernel satisfies $\Lambda_{t,t+1}^{j,h} \geq 1$ for liquid securities (those entering L_t^h with a positive partial derivative). It *augments* the standard SDF, effectively reducing the discount rate for liquid assets and generating a *liquidity premium*.

Iterating forward, the security price decomposes into three components:

$$p_t^j = \underbrace{\lim_{T\to\infty} \mathbb{E}_t\left[\sum_{s=1}^{T} \mu_{t,t+s}^h z_{t+s}^j\right]}_{\text{Fundamental value } F_t^j} + \underbrace{\lim_{T\to\infty} \mathbb{E}_t\left[\sum_{s=1}^{T} \mu_{t,t+s}^h (\Lambda_{t,t+s}^{j,h} - 1) z_{t+s}^j\right]}_{\text{Liquidity premium } LP_t^j} + \underbrace{\lim_{T\to\infty} \mathbb{E}_t\left[\mu_{t,t+T}^h \Lambda_{t,t+T}^{j,h} p_{t+T}^j\right]}_{\text{Bubble component } B_t^j} \tag{13}$$

where $\Lambda_{t,t+s}^{j,h} = \prod_{\tau=t}^{t+s-1} \Lambda_{\tau,\tau+1}^{j,h}$ is the multi-period liquidity kernel.

This decomposition is central to my theory. The *fundamental value* F_t^j reflects the asset's cash flows discounted at the standard rate. The *liquidity premium* LP_t^j reflects the additional value from the asset's role in relaxing liquidity constraints. The *bubble component* B_t^j reflects the residual value from indefinite resalability. In a crisis, as I will show below, it is the

collapse of LP_t^j—not F_t^j—that triggers the discontinuous contraction of liquid media.

2.7 Regularity Condition

The liquidity kernel $\Lambda_{t,t+1}^{j,h}$ is well-defined and strictly positive if and only if the denominator in Equation (11) is strictly positive, that is,

$$\lambda_t^h \cdot \frac{\partial L_t^h}{\partial (a_t^{j,h} p_t^j)} < 1. \tag{14}$$

This restriction guarantees that the pricing operator remains finite and that marginal valuations are well-behaved.

Define

$$v_t^j \equiv \frac{\partial L_t^h}{\partial (a_t^{j,h} p_t^j)}$$

as the asset's *liquidity weight*. Thus, the parameter v_t^j measures the marginal expansion of transaction capacity generated by one additional numeraire unit invested in security j. In other words, v_t^j captures the *moneyness* of asset j at time t. Assets with high transactional acceptability have v_t^j close to one, whereas fundamentally illiquid backing assets have $v_t^j = 0$. Using this notation, (14) becomes

$$\lambda_t^h v_t^j < 1.$$

The term $\lambda_t^h v_t^j$ represents the liquidity-adjusted marginal valuation of wealth invested in asset j. The regularity condition therefore requires that the marginal utility gain from liquidity services remain strictly below the marginal utility of wealth itself. When $\lambda_t^h v_t^j \to 1$, the denominator of the liquidity kernel collapses. The kernel diverges, the associated asset price is bid upward without bound, and the equilibrium ceases to be well-defined. This breakdown reflects an unbounded marginal willingness to pay for liquidity rather than an immediate disappearance of fundamentals.

Within the monetary production framework, the mathematical singularity of the pricing operator defines the absolute upper bound of the system. However, the economically relevant crisis boundary, $\tilde{q}^*$, is reached strictly before this singular point. Thus, the breakdown of equilibrium is governed not by the formal divergence of the operator, but by the earlier violation of the solvency condition that sustains monetary circulation.

A decline in the fundamental quality of backing assets contracts the real resources supporting liquid media, tightening liquidity constraints and driving up the shadow value λ_t^h. As this deterioration pushes $\lambda_t^h v_t^j$ toward its critical limit of 1, the system approaches a market freeze. I formally define this measure of backing-asset quality as $\tilde{q}_t$ in the following

section, and derive its exact critical threshold, $\tilde{q}^*$, in Section 3.

2.8 Tractable Specification

To derive closed-form results, I specialize the environment to the following economy.

Assumption 2 (Cash-in-Advance and Logarithmic Utility)**.**

(i) **Preferences.** Households have time-separable log utility over a single consumption good:

$$U^h = \mathbb{E}_0\left[\sum_{t=0}^{\infty} \beta^t \ln(c_t^h)\right] \tag{15}$$

where $\beta \in (0,1)$ is the common discount factor.

(ii) **Liquidity constraint.** The economy uses a cash-in-advance constraint where consumption c_t^h (priced in numeraire) must be backed by liquid holdings:

$$L_t^h = a_t^{M,h} p_t^M + \sum_b a_t^{D^b,h} p_t^{D^b} v_t^{D^b} - c_t^h \geq 0 \tag{16}$$

where $v_t^{D^b} \in [0,1]$ is the liquidity weight of bank b's deposits. When deposits are fully accepted as money, $v_t^{D^b} = 1$ and deposits are a perfect substitute for outside money. When the bank's solvency is in doubt, $v_t^{D^b}$ falls.

(iii) **Technology.** The consumption good is produced with a constant endowment $Y > 0$.

(iv) **Banking.** A representative bank holds backing assets A_t and issues deposits D_t plus equity E_t:

$$A_t = D_t + E_t \tag{17}$$

Backing assets are perpetuities yielding a constant real cash flow y^A per period.

Under this specification, the liquidity kernel simplifies. Because the CIA (16) depends only on current-period portfolio positions $a_t^{j,h} p_t^j$ (not on inherited positions), the household's first-order condition for deposits of a solvent bank with a constant liquidity weight $v_t^{D^b} = v$ yields the kernel (see Section A.1 for the derivation):

$$\Lambda = \frac{1}{1 - \lambda_t v} \tag{18}$$

From the household's first-order condition for outside money (with $v^M = 1$) in steady state, the shadow price of liquidity is:

$$\lambda = \frac{i}{1+i} \tag{19}$$

where i is the nominal interest rate. In steady state ($\lambda_t = \lambda$), the kernel becomes:

$$\Lambda = \frac{1}{1 - \lambda v} \tag{20}$$

The key parameter governing this multiplier is $\omega \equiv \lambda v$, which I define as the *effective liquidity intensity* of the economy. When ω is small (low nominal rates, low velocity weight), the kernel is close to 1 and money is nearly superneutral. When ω is large, the kernel amplifies all asset price movements, pushing the economy closer to the crisis boundary.

Backing-asset quality. I formally define the quality of the banking sector's asset portfolio as an *exogenous fundamental state variable*:

$$\tilde{q}_t \in (0, 1] \tag{21}$$

The parameter $\tilde{q}_t$ determines the fraction of promised cash flows that the backing assets actually deliver. A fully performing asset has $\tilde{q}_t = 1$; an impaired asset has $\tilde{q}_t < 1$. The fundamental value of the backing-asset portfolio is therefore:

$$F_t = \frac{\tilde{q}_t \cdot y^A}{r} \tag{22}$$

where $r = 1/\beta - 1$ is the steady-state real interest rate, and y^A is the promised (unimpaired) cash flow. Quality $\tilde{q}_t$ is *exogenous to the monetary equilibrium*: it is determined by real-sector fundamentals (borrower solvency, asset performance, interest-rate shocks to duration) that are independent of whether the bank's liabilities circulate as money. This fundamental separability is essential for the crisis mechanism: it ensures that the self-referential loop operates through the liquidity premium alone, while the quality shock provides the exogenous impulse that triggers the equilibrium transition.

As a diagnostic, the endogenous ratio F_t/A_t reveals the fraction of the bank's market value attributable to fundamentals versus the liquidity premium. At the solvency boundary $E_t = 0$, this ratio coincides with $\tilde{q}_t \cdot y^A/(r \cdot A_t)$, providing a bridge between the exogenous quality parameter and observable balance-sheet ratios.

The tractable specification is a special case of the general framework with specific functional forms. All results derived below hold exactly in this specification and extend (with appropriate continuity and monotonicity conditions) to the broader class of liquidity constraints satisfying the regularity condition (14). The CIA structure is not essential; what is essential is that some assets enter liquidity constraints, that banks issue liabilities that do so, and that the solvency of banks depends on the quality of their backing assets. These features are present

in any specification of L_t^h that accommodates inside money.

The cash-in-advance specification is chosen for analytical tractability. It generates a binding constraint with a well-defined shadow price λ_t, which maps directly to the crisis threshold. Under money-in-utility, the liquidity motive operates through a first-order condition rather than a binding constraint, yielding qualitatively identical results but without the sharp threshold. I discuss MIU as a robustness check in Section A.5.

3 The Monetary Production Technology

The preceding analysis treats the liquidity constraint L_t^h as a restriction on household optimization. In this section, I pivot from the demand for liquidity to its supply by formalizing the *monetary production technology* that relaxes this constraint. Rather than viewing liquidity merely as an exogenous property of assets, I model it as the endogenous output of a balance-sheet transformation: a mapping from a fundamentally less liquid portfolio of backing assets into liquid transaction media.

In this framework, backing assets constitute inputs, liquid media constitute outputs, and the liquidity kernel prices the transformation. The shadow value λ_t thus acquires a dual interpretation: it represents both the household's marginal value of relaxing the liquidity constraint and the marginal product of liquidity in the production technology.

Under this perspective, solvency conditions translate directly into feasibility constraints for liquidity creation. The quality of backing assets, $\tilde{q}_t$, determines the effective input available to the technology. A deterioration in $\tilde{q}_t$ reduces productive capacity, tightening equilibrium constraints and raising λ_t. The crisis boundary characterized in Section 2.7 thus emerges as a technological breakdown: when input quality falls below the threshold $\tilde{q}^*$, the liquidity production function can no longer sustain a finite equilibrium price for liquid liabilities.

Ultimately, this approach endogenizes the liquidity premium as the return to a balance-sheet transformation, strictly linking monetary instability to the fundamental quality of backing assets rather than to exogenous shifts in beliefs.

3.1 Definitions

Definition 3 (Monetary Production Technology)**.** A **monetary production technology** is a mapping

$$\mathcal{T}_t : \left\{a_t^{j,i}\, p_t^j\right\}_{j\in\mathcal{J}_B} \longrightarrow \ell_t \in \mathbb{R}_+$$

that transforms a portfolio of **backing assets** $\{a_t^{j,i}\, p_t^j\}_{j\in\mathcal{J}_B}$, where $\mathcal{J}_B \subset \mathcal{J}$ denotes the subset of securities held on the producer's balance sheet, into a flow of **liquidity services** ℓ_t.

This flow is defined as the maximum transaction volume that can be sustained by the liquid media supplied against those assets:

$$\ell_t \equiv \max_{x_t^{1,h}} x_t^{1,h} \quad \text{subject to} \quad L_t^h\Big(\{a_t^{j,h} p_t^j\}_j, \dots\Big) \geq 0 \tag{23}$$

The technology is parameterized by **backing-asset quality**:

$$\tilde{q}_t \in (0, 1] \tag{24}$$

an exogenous fundamental state variable that determines the fraction of promised cash flows actually delivered by the backing assets. The fundamental value of the backing-asset portfolio is $F_t = \tilde{q}_t \cdot y^A / r$, where y^A is the promised (unimpaired) cash flow and $r = 1/\beta - 1$ is the steady-state real interest rate.

The definition captures a simple but powerful idea: money does not emerge from nothing. Every form of money—whether bank deposits, fiat currency, stablecoins, or even minted commodity money—is a liquid medium derived from some underlying asset or portfolio. The fundamental quality of those assets determines the viability of the money.

The quality parameter $\tilde{q}_t$ is *exogenous to the monetary equilibrium*—it is determined by real-sector fundamentals (borrower solvency, asset performance, duration risk) that are independent of whether the liquid media circulate as money. This fundamental separability is the key structural assumption. The market value of backing assets $A_t = F_t + LP_t$ is, by contrast, an equilibrium object: the liquidity premium LP_t is present if and only if the liquid media are accepted, and its magnitude depends on the monetary environment. The self-referential loop runs strictly through LP_t—not through $\tilde{q}_t$. When I refer to "fundamental separability" in condition (A3) of the Unification Proposition, I mean exactly this: $\tilde{q}_t$ is separable from the monetary equilibrium. For the core model of Section 3, this separability is a direct consequence of the Cash-in-Advance specification (the backing-asset cash flow y^A is a real endowment, independent of whether deposits circulate). For the unification framework of Section 4, separability is a *sufficient condition* imposed on each arrangement—one that can fail when backing assets are themselves monetary instruments.

Definition 4 (Monetary Producer)**.** An agent i (bank, central bank, government, mint, or non-bank financial institution) is a **monetary producer** if it supplies assets j that enter the liquidity constraints of households with a strictly positive liquidity weight $v_t^j > 0$ for some household h.

The **production margin** of producer i is the spread between the cash flows generated

by its backing assets and the cost of servicing its liquid liabilities:

$$\pi_t^i = \sum_{j \in \mathcal{J}_B} a_t^{j,i} \, z_t^j - \sum_{k \in \mathcal{J}_L} a_t^{k,i} \, z_t^k \tag{25}$$

where $\mathcal{J}_B$ is the set of backing assets held, and $\mathcal{J}_L$ is the set of liquid media issued by i (where $a_t^{k,i} > 0$ denotes the quantity issued).

The monetary production margin π_t^i naturally nests several familiar concepts from the literature:

(i) *Bank franchise value* (Calomiris and Kahn, 1991): the production margin π_t^i is positive because banks finance high-fundamental-value assets with low-fundamental-value (but high-liquidity-premium) liabilities.

(ii) *Seigniorage*: for a government issuing fiat currency, π_t^i is the seigniorage revenue—the return on assets acquired by issuing zero-cost liabilities.

(iii) *Money creation* (McLeay et al., 2014): the act of a bank making a loan and simultaneously creating a deposit is precisely the operation of $\mathcal{T}_t$—transforming a loan (backing asset) into a deposit (liquid medium).

3.2 The Production Frontier

I now characterize the set of feasible (backing-asset quality, liquidity output) pairs. To derive this production frontier in closed form, I specialize the general technology back to the tractable Cash-in-Advance economy introduced in Section 2.8.

Consider a representative monetary producer that holds backing assets A_t and issues liquid media $\mathcal{M}_t$ (deposits, notes, coins, or any instrument entering L_t^h) plus equity E_t:

$$A_t = \mathcal{M}_t + E_t \tag{26}$$

The market value of backing assets decomposes, following (13), into fundamental value and liquidity premium (setting the bubble component to zero for backing assets):

$$A_t = F_t + LP_t \tag{27}$$

where $F_t = \tilde{q}_t \cdot y^A / r$ is the fundamental value. The quality parameter $\tilde{q}_t \in (0, 1]$ is exogenous to the producer—it represents the fraction of promised cash flows that the backing assets actually deliver.

The backing assets carry a liquidity premium LP_t because they support the issuance of liquid media. From the steady-state Euler equation (12) with discount factor $\mu = \beta$ and liquidity kernel $\Lambda = 1/(1 - \lambda v)$, the price of a perpetual asset with cash flow z is:

$$p = \frac{\beta \Lambda z}{1 - \beta \Lambda} \tag{28}$$

and the fundamental value (with $\Lambda = 1$, i.e., $\lambda = 0$) is $F = \beta z/(1 - \beta) = z/r$. The liquidity premium is therefore:

$$LP = p - F = \frac{\beta z(\Lambda - 1)}{(1 - \beta\Lambda)(1 - \beta)} \tag{29}$$

Substituting $\Lambda - 1 = \omega/(1 - \omega)$ and $1 - \beta\Lambda = (1 - \omega - \beta)/(1 - \omega)$, and defining $\omega \equiv \lambda v$ as the *effective liquidity intensity*, the liquidity premium can be written as:

$$LP = \Phi(\omega, \beta) \cdot z \tag{30}$$

where z is the real cash flow of the backing assets and $\Phi(\omega, \beta)$ is the **liquidity production coefficient**:

$$\Phi(\omega, \beta) \equiv \frac{\beta\omega}{(1 - \omega - \beta)(1 - \beta)} \tag{31}$$

For a backing asset with quality $\tilde{q}_t$ delivering cash flow $z = \tilde{q}_t \cdot y^A$, the liquidity premium is $LP = \Phi \cdot \tilde{q}_t \cdot y^A$.

To connect this balance-sheet capacity to the real economy, I impose money market clearing. In equilibrium, with the Cash-in-Advance constraint binding, the real supply of liquid media must equal aggregate output:

$$\frac{\mathcal{M}_t}{P_t} = Y_t \tag{32}$$

Assumption 5 (Parameter Restriction)**.** The effective liquidity intensity satisfies $\omega < 1 - \beta$, ensuring that the liquidity production coefficient $\Phi(\omega, \beta)$ is positive, the liquidity kernel $\Lambda = 1/(1 - \omega)$ is bounded, and the steady-state asset price $p = \beta\Lambda z/(1 - \beta\Lambda)$ is well-defined.

Before characterizing the frontier, it is useful to formally distinguish between liquidity capacity and realized transactions. The object ℓ_t denotes *liquidity capacity*: the maximum transaction volume the producer's balance sheet can technologically support given its asset composition and quality.

By contrast, actual transactions are determined in equilibrium by goods-market clearing. Under the Cash-in-Advance condition, real balances satisfy $\mathcal{M}_t/P_t = Y_t$, so that realized transaction volume equals aggregate output. Thus, ℓ_t and Y_t play conceptually distinct roles:

the former is a supply-side feasibility constraint, whereas the latter is a general equilibrium outcome pinned down by real activity.

The frontier $\mathcal{F}(\tilde{q}_t)$ operationalizes this distinction. For a given level of backing-asset quality $\tilde{q}_t$, the frontier determines the maximum stock of liquid media that can be credibly sustained by the producer. Formally, $\mathcal{F}(\tilde{q}_t)$ is the largest $\mathcal{M}_t/P_t$ consistent with both balance-sheet solvency ($E_t \geq 0$) and the regularity condition in Section 2.7 that ensures the liquidity kernel Λ remains bounded.

If $\mathcal{F}(\tilde{q}_t) \geq Y_t$, the producer can sustain the equilibrium money supply required for goods-market clearing while maintaining positive equity. In this region, liquidity capacity exceeds realized transactions, making monetary production feasible. If instead $\mathcal{F}(\tilde{q}_t) < Y_t$, the balance sheet cannot support the required real balances without violating solvency or regularity. Monetary production therefore becomes technologically infeasible.

Proposition 6 (Production Frontier)**.** *Under the CIA specification with log utility and Theorem 5, the monetary production frontier $\mathcal{F}(\tilde{q}_t) = \max \ell_t$ subject to the producer's solvency constraint $E_t \geq 0$ satisfies:*

(i) $\mathcal{F}$ is strictly increasing in $\tilde{q}_t$: higher backing-asset quality sustains greater liquidity output.

(ii) $\mathcal{F}$ is linear in $\tilde{q}_t$: each unit of quality improvement generates the same increment of liquidity capacity.

(iii) There exists a critical threshold $\tilde{q}^* = (1-\beta-\omega)/(\beta\kappa) > 0$ such that $\mathcal{F}(\tilde{q}_t) = 0$ for all $\tilde{q}_t \leq \tilde{q}^*$: below the threshold, no positive level of monetary production can be sustained in equilibrium. When $\kappa > (1-\beta-\omega)/\beta$, the threshold satisfies $\tilde{q}^* < 1$ and a feasible monetary equilibrium exists for sufficiently high asset quality. When $\kappa \leq (1-\beta-\omega)/\beta$, the threshold satisfies $\tilde{q}^* \geq 1$: the fundamental cash-flow base is too small for monetary production to be feasible at any asset quality—the economy is in a regime of structural monetary infeasibility.[2]

Proof. (i)–(ii). At the solvency boundary $E_t = 0$, the balance sheet gives $A_t = \mathcal{M}_t$, i.e., $F_t + LP_t = \mathcal{M}_t$. Substituting $F_t = \tilde{q}_t \cdot y^A/r$ and $LP_t = \Phi \cdot \tilde{q}_t \cdot y^A$ (from (30) with $z = \tilde{q}_t \cdot y^A$), and using money market clearing $\mathcal{M}_t/P_t = Y_t$, the maximum real liquidity output sustainable at quality $\tilde{q}_t$ is:

$$\mathcal{F}(\tilde{q}_t) = \tilde{q}_t \cdot y^A \left(\frac{1}{r} + \Phi\right) = \tilde{q}_t \cdot \kappa \cdot Y_t \cdot \frac{\beta}{1-\omega-\beta}$$

[2]The condition $\tilde{q}^* \geq 1$ is the extreme form of the permanent fragility discussed in Section 3.6.2. While $\tilde{q}^{**} > 1$ means the safe region is unreachable, $\tilde{q}^* \geq 1$ means even the fragile monetary equilibrium cannot exist.

where $\kappa \equiv y^A/Y_t$. This is strictly increasing and linear in $\tilde{q}_t$, with slope $\kappa Y_t \beta/(1-\omega-\beta) > 0$.

(iii). The frontier reaches the CIA clearing level Y_t at the critical quality:

$$\mathcal{F}(\tilde{q}^*) = Y_t \implies \tilde{q}^* \cdot \kappa \cdot \frac{\beta}{1-\omega-\beta} = 1$$

Solving:

$$\boxed{\tilde{q}^* = \frac{1-\beta-\omega}{\beta\kappa}} \tag{33}$$

For $\tilde{q}_t < \tilde{q}^*$, the solvency constraint is violated for any $m_t > 0$ that satisfies money market clearing. Hence $\mathcal{F}(\tilde{q}_t) = 0$: the monetary equilibrium cannot exist. □

The linearity of the production frontier has a precise economic interpretation: within the monetary equilibrium, each unit of quality improvement generates the same increment of liquidity capacity. The *abruptness* of financial crises does not arise from the frontier's curvature—it arises from the *equilibrium structure*. At $\tilde{q}^*$, the self-sustaining monetary equilibrium ceases to exist (a saddle-node bifurcation), and the economy discontinuously transitions to the non-monetary equilibrium. The frontier is a smooth, linear function of quality, but the equilibrium selection is a step function at the threshold. This is why financial crises appear sudden despite gradual underlying deterioration: the fundamental shock is continuous, but the equilibrium response is discontinuous.

The threshold $\tilde{q}^*$ depends on three groups of primitives: (i) the **monetary environment** ($\omega = \lambda v$); (ii) the **real economy** ($\kappa = y^A/Y$); and (iii) **time preference** (β)[3]. Crucially, $\tilde{q}^*$ does not depend on the *identity* of the backing assets or the *institutional form* of money.

A bank issuing deposits against mortgages, a currency board holding foreign reserves, and a government issuing fiat backed by tax revenue all face the same production-technology constraint. As will be explored in Section 3.3, once household expectations are endogenous, a second threshold $\tilde{q}^{**}$ emerges. This defines the region of fragility $\mathcal{G}$ where the system is susceptible to sunspot-driven, self-fulfilling crises (see Figure 1).

3.3 The Crisis Theorem

The production frontier establishes the minimum technological requirement for monetary existence. I now characterize the full equilibrium structure, showing how the interaction between balance-sheet solvency and household expectations generates a zone of endogenous fragility.

[3]The discount factor β uniquely determines the steady-state real interest rate via $r = 1/\beta - 1$.

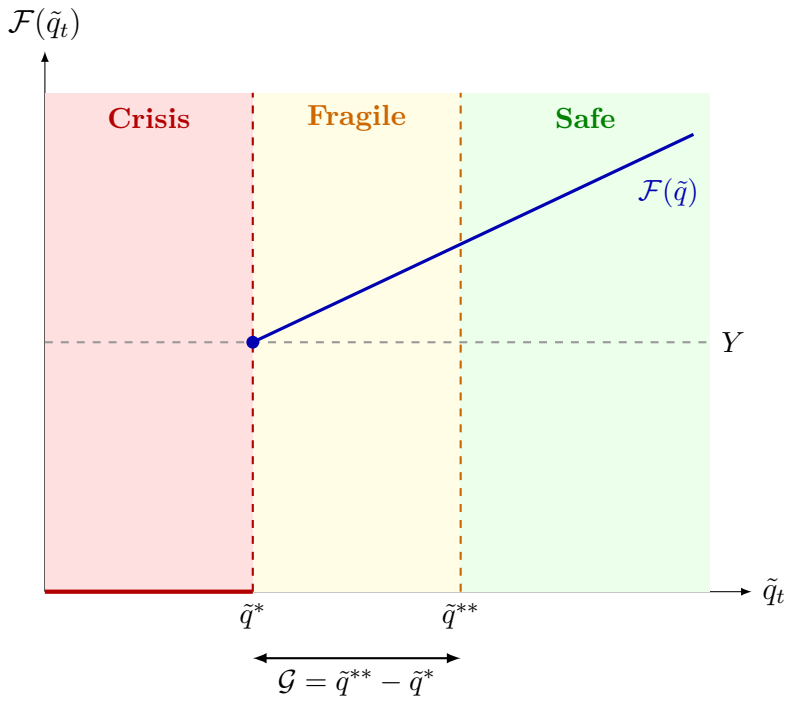


Figure 1: The monetary production frontier. $\mathcal{F}(\tilde{q})$ is the maximum stock of liquid media sustainable at backing-asset quality $\tilde{q}$. The frontier is zero for $\tilde{q} \leq \tilde{q}^*$ (crisis region), increasing and linear for $\tilde{q} > \tilde{q}^*$. The horizontal dashed line Y marks the CIA clearing level: $\mathcal{F}(\tilde{q}) \geq Y$ (positive equity) requires $\tilde{q} \geq \tilde{q}^*$. The fragility index $\mathcal{G} = \tilde{q}^{**} - \tilde{q}^*$ measures the width of the fragile region where sunspot-driven crises are possible.

Theorem 7 (Crisis Threshold)**.** *In the CIA economy with log utility, consider a representative monetary producer issuing liquid media $\mathcal{M}_t$ backed by assets of quality $\tilde{q}_t$. There exist critical thresholds $\tilde{q}^*$ and $\tilde{q}^{**}$ (with $\tilde{q}^* < \tilde{q}^{**}$) such that:*

(i) ***Safe region*** *($\tilde{q}_t \geq \tilde{q}^{**}$): There exists a unique equilibrium with $\mathcal{M}_t > 0$. The producer is solvent based on the fundamental value of its assets alone; no liquidity premium is required to maintain a positive equity buffer.*

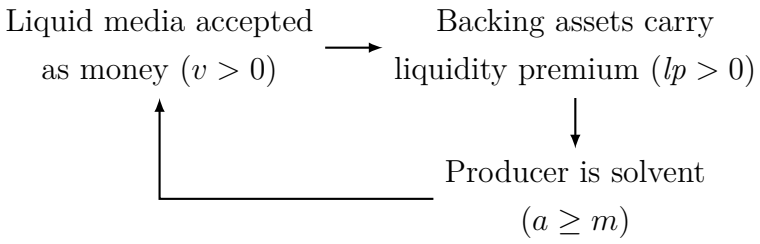
(ii) ***Fragile region*** *($\tilde{q}^* < \tilde{q}_t < \tilde{q}^{**}$): Two equilibria coexist:*

- *A monetary equilibrium with $\mathcal{M}_t > 0$, in which backing assets carry a positive liquidity premium that sustains the producer's solvency.*
- *A non-monetary equilibrium with $\mathcal{M}_t = 0$, in which the liquid media lose their acceptability, the liquidity premium collapses, and monetary exchange breaks down.*

(iii) ***Crisis region*** *($\tilde{q}_t \leq \tilde{q}^*$): The unique equilibrium is the non-monetary equilibrium ($\mathcal{M}_t = 0$). Monetary production is impossible because even the maximum liquidity premium supported by the technology cannot restore the producer's solvency.*

Proof. (i). When $LP = 0$ (no liquidity premium), the market value of assets is purely fundamental: $A = F = \tilde{q} \cdot y^A/r$. Real solvency without the liquidity premium requires $F \geq m$, where $m = Y$ from the CIA constraint. This implies $\tilde{q} \cdot y^A/r \geq Y$, which yields $\tilde{q} \geq r/\kappa \equiv \tilde{q}^{**}$. For $\tilde{q} \geq \tilde{q}^{**}$, the producer is solvent even at $LP = 0$. The monetary equilibrium is unique: a non-monetary equilibrium requires $LP = 0$, but with $LP = 0$ the producer remains fundamentally solvent. Therefore, its liquid media remain safe, carry a positive v_t, and thus generate a positive liquidity premium, contradicting the premise that $LP = 0$.

(ii). For $\tilde{q}^* < \tilde{q} < \tilde{q}^{**}$, the producer requires a positive liquidity premium for solvency: evaluated in real terms, $f + lp \geq m$ holds only if $lp > 0$. The liquidity premium is strictly positive if and only if the liquid media are accepted as money (i.e., $v > 0$ in the CIA constraint), which in turn requires the producer to be solvent. This creates a self-referential fixed point:



Both the self-sustaining loop (monetary equilibrium) and its absence (non-monetary equilibrium) are consistent with rational expectations. The economy is in the *fragile region* where sunspot-driven crises—self-fulfilling losses of confidence—are mathematically possible.

(iii). For $\tilde{q} < \tilde{q}^*$, even the self-sustaining loop cannot close. From Theorem 6(iii), the maximum real liquidity premium (computed under the assumption that the liquid media are fully accepted as money) is technologically insufficient: $f + lp^{\max} < m$. The monetary equilibrium cannot exist. The unique equilibrium is the non-monetary equilibrium ($m = 0$).

The transition at $\tilde{q} = \tilde{q}^*$ represents a critical regime shift. As $\tilde{q}$ approaches $\tilde{q}^*$ from above, the monetary equilibrium and the boundary of the feasible solvency region collide. For $\tilde{q} < \tilde{q}^*$, the solvency constraint is strictly violated for any positive level of real balances. Crucially, below $\tilde{q}^*$, a monetary collapse is not a coordination failure (as in Diamond and Dybvig, 1983) but a fundamental impossibility. □ □

Remark 8 (Relationship to Brunnermeier-Niepelt). This result is not contained in Brunnermeier and Niepelt (2019). BN characterize properties of the monetary equilibrium *conditional on its existence*: they derive conditions under which private and public money produce equivalent allocations. They do not ask when the monetary equilibrium fails to exist, derive

the quality threshold below which failure is inevitable, or characterize the three-region equilibrium structure. The production-technology reinterpretation is not cosmetic: it reframes the liquidity constraint as a feasibility boundary, revealing $\tilde{q}^*$ as the point at which the production set collapses. The closed-form expression $\tilde{q}^* = (1 - \beta - \omega)/(\beta\kappa)$ and the comparative statics of Theorem 12 are new.

Crucially, the collapse ($\mathcal{M}_t = 0$) characterizes the failure of the *specific liquid media* issued by the failing producer. While the formal model represents this limit as a non-monetary equilibrium, in practice, an economy typically falls back on alternative, less efficient monetary instruments rather than reverting to pure barter. In a banking crisis, the economy reverts to outside money alone (e.g., physical cash); in a currency crisis, the economy dollarizes, switching to a foreign currency. In all these cases, the effective shadow cost of liquidity λ_t spikes because the fallback instruments cannot efficiently support the required transaction volume. The real resource cost of exchange rises sharply, and the resulting welfare loss—the difference in utility between the monetary and non-monetary equilibria—can be directly evaluated from the model's primitives.

3.4 Amplification

The crisis threshold $\tilde{q}^*$ determines the regions of the state space. Within those regions, the model generates severe amplification dynamics as the system approaches the boundary.

Proposition 9 (Amplification Mechanism)**.** *When backing-asset quality $\tilde{q}_t$ declines (but remains above $\tilde{q}^*$), the following feedback loop operates:*

(i) Lower $\tilde{q}_t$ reduces the producer's real equity buffer: $e_t = f_t + lp_t - m_t$ shrinks.

(ii) Reduced equity increases the producer's fragility: the distance to the critical threshold $\tilde{q}^*$ narrows.

(iii) If the solvency constraint binds, the contraction in real liquidity capacity forces a reduction in real balances m_t. This tightens the CIA constraint, raising the shadow cost of liquidity λ_t through goods-market clearing.

(iv) Higher λ_t increases the liquidity kernel Λ_t, amplifying the sensitivity of asset prices to liquidity conditions.

(v) If the producer is forced to liquidate backing assets to meet redemptions (fire sales), aggregate liquidation depresses asset prices further, accelerating the collapse of the liquidity premium independent of the fundamental state $\tilde{q}_t$.

(vi) The feedback continues until either a new interior monetary equilibrium is reached (partial contraction) or $\tilde{q}_t$ crosses $\tilde{q}^*$ (full collapse).

Proof. Parts (i)–(ii) follow directly from the real balance-sheet identity and the definition of $\tilde{q}^*$.

Part (iii): At the solvency boundary, real balances are constrained by capacity: $m_t = f_t + lp_t$. A decline in $\tilde{q}_t$ (holding lp_t constant) reduces fundamental value f_t and hence forces a contraction in m_t. By the CIA constraint, a lower m_t restricts real transaction volume; this scarcity of liquidity strictly raises its shadow cost λ_t in equilibrium.

Part (iv): Evaluating the liquidity kernel $\Lambda_t = 1/(1 - \lambda_t v_t)$, the partial derivative with respect to the current shadow cost is:

$$\frac{\partial \Lambda_t}{\partial \lambda_t} = \frac{v_t}{(1 - \lambda_t v_t)^2} > 0$$

This derivative is strictly increasing and convex in λ_t, demonstrating that amplification accelerates non-linearly as the system approaches the crisis boundary.

Part (v) appeals to the standard fire-sale externality widely established in the macro-finance literature (e.g., Brunnermeier and Sannikov, 2014), where aggregate liquidation depresses asset prices.

Part (vi) follows from the fixed-point structure defined in Theorem 7: the monetary equilibrium is locally stable for $\tilde{q} > \tilde{q}^*$ but definitively ceases to exist at $\tilde{q} = \tilde{q}^*$. □ □

Remark 10 (Endogenous velocity). In the baseline model, the velocity weight v is fixed. An extension that allows the velocity weight to respond to the producer's solvency—for example, $v_t = v \cdot \mathbf{1}[e_t/a_t > \bar{e}]$ for some minimum real equity ratio $\bar{e} \geq 0$—would introduce an additional amplification channel: declining solvency reduces the money-ness of the liquid media directly. The case $\bar{e} = 0$ recovers the baseline model; the case $\bar{e} > 0$ captures the empirical pattern in which depositor unease and velocity contractions begin prior to formal insolvency. I leave the formal analysis of endogenous velocity to future work.

3.5 Comparative Statics of the Crisis Threshold

The closed-form expression (33) yields testable comparative statics.

Definition 11 (Financial Fragility)**.** Define the **fragility index** as

$$\mathcal{G}(\omega, \kappa, \beta) \equiv \tilde{q}^{**} - \tilde{q}^* = \frac{\omega}{\beta\kappa} \tag{34}$$

Fragility $\mathcal{G}$ measures the width of the fragile region—the set of backing-asset qualities for which a sunspot-driven crisis is possible but not inevitable. When $\mathcal{G} > 0$, the three-region structure exists; when $\mathcal{G} \leq 0$, the economy transitions directly from safe to crisis (knife-edge case). The fragility index is directly proportional to the effective liquidity intensity ω and inversely proportional to the discounted cash-flow base $\beta\kappa$.

Corollary 12 (Determinants of Financial Fragility)**.** *The crisis threshold $\tilde{q}^*$ and the fragility index $\mathcal{G}$ satisfy:*

(i) $\partial\tilde{q}^*/\partial\omega < 0$ *and* $\partial\mathcal{G}/\partial\omega > 0$*: higher effective liquidity intensity (higher nominal interest rates or higher velocity) lowers the crisis threshold but widens the fragile region. Moreover, the amplification mechanism (Theorem 9) is stronger at high ω because $\partial\Lambda_t/\partial\lambda_t$ is increasing and convex. The net effect is that economies with higher ω are more fragile: the fragile region is wider, amplification within it is stronger, and the susceptibility to crossing $\tilde{q}^*$ increases.*

(ii) $\partial\tilde{q}^*/\partial\kappa < 0$ *and* $\partial\mathcal{G}/\partial\kappa < 0$*: economies whose backing assets generate a larger share of output have both a lower crisis threshold and a narrower fragile region, making them unambiguously more resilient. Intuitively, a larger fundamental asset base κ pushes both $\tilde{q}^*$ and $\tilde{q}^{**}$ downward, but the fragile gap $\mathcal{G} = \omega/(\beta\kappa)$ also shrinks because the monetary system's reliance on the endogenous liquidity premium diminishes.*

(iii) $\partial\tilde{q}^*/\partial\beta < 0$*: more patient economies are unambiguously more resilient (lower $\tilde{q}^*$). Higher β raises steady-state asset prices and the equilibrium liquidity premium, increasing the monetary system's capacity to absorb fundamental shocks without breaching the solvency constraint.*

Proof. From $\tilde{q}^* = (1-\beta-\omega)/(\beta\kappa)$, $\tilde{q}^{**} = (1-\beta)/(\beta\kappa)$, and $\mathcal{G} = \omega/(\beta\kappa)$:

(i). $\partial\tilde{q}^*/\partial\omega = -1/(\beta\kappa) < 0$; $\partial\mathcal{G}/\partial\omega = 1/(\beta\kappa) > 0$ (see Section A.4).

(ii). $\partial\tilde{q}^*/\partial\kappa = -(1-\beta-\omega)/(\beta\kappa^2) = -\tilde{q}^*/\kappa < 0$; $\partial\tilde{q}^{**}/\partial\kappa = -(1-\beta)/(\beta\kappa^2) < 0$; $\partial\mathcal{G}/\partial\kappa = -\omega/(\beta\kappa^2) < 0$. All three effects have unambiguous signs: higher κ lowers the crisis threshold, the safe threshold, and the width of the fragile region.

(iii). $\partial\tilde{q}^*/\partial\beta = -(1-\omega)/(\beta^2\kappa) < 0$, since $\omega < 1-\beta < 1$ by the parameter restriction (Theorem 5). □ □

These comparative statics generate sharp cross-country predictions. Economies with high nominal interest rates or high monetary velocity (high ω) should feature wider fragile regions and stronger amplification, making them structurally more crisis-prone. Economies with a larger fundamental asset base eligible to back money relative to GDP (high κ) have

both lower crisis thresholds and narrower fragile regions, making them unambiguously more resilient to fundamental shocks. Finally, economies with high discount rates (low β) face higher thresholds, making monetary production inherently more precarious. Notably, all three comparative statics have *unambiguous* signs under the simplified CIA kernel—a sharpening relative to more general specifications where some effects are ambiguous (see Section A.4).

3.6 Quantitative Illustration

To establish that the crisis threshold $\tilde{q}^*$ lies in an empirically relevant range, I evaluate it at standard parameter values. The exercise is illustrative, not a full structural estimation: its purpose is to demonstrate the three-region equilibrium structure, verify that the comparative statics operate in the predicted directions, and show that the model's classification of safe, fragile, and crisis institutions maps onto observed outcomes.

3.6.1 Idiosyncratic Avoidable Fragility: Baseline and SVB Example

I set baseline annual parameters: $\beta = 0.96$ (real rate $r = 1/\beta - 1 \approx 4.17\%$), nominal interest rate $i = 0.01$ ($\lambda = i/(1+i) \approx 0.0099$), velocity weight $v = 0.5$ (liquid media cover half of transaction needs), and $\kappa = y^A/Y = 0.048$ (backing-asset cash flows equal 4.8% of GDP). The effective liquidity intensity is $\omega = \lambda v \approx 0.00495$.[4]

From these primitives, the model yields:

$$\Lambda = \frac{1}{1-\omega} = \frac{1}{0.99505} \approx 1.005 \tag{35}$$

$$\Phi(\omega, \beta) = \frac{\beta\omega}{(1-\omega-\beta)(1-\beta)} \approx 3.39 \tag{36}$$

$$\tilde{q}^* = \frac{1-\beta-\omega}{\beta\kappa} = \frac{0.03505}{0.04608} \approx \mathbf{0.761} \tag{37}$$

$$\tilde{q}^{**} = \frac{1-\beta}{\beta\kappa} = \frac{0.04}{0.04608} \approx \mathbf{0.868} \tag{38}$$

A crisis threshold of $\tilde{q}^* \approx 0.761$ means that monetary production collapses when backing assets deliver only 76 cents of fundamental value per dollar of promised cash flows. The fragile region spans $(0.761, 0.868)$: for backing-asset quality in this range, the monetary equilibrium exists but a sunspot-driven crisis is possible. The fragility index is $\mathcal{G} = \tilde{q}^{**} - \tilde{q}^* = \omega/(\beta\kappa) \approx$

[4]The parameter restriction (Theorem 5) requires $\omega < 1 - \beta = 0.04$ for the perpetuity-based asset pricing formula to converge ($\beta\Lambda < 1$). This bound arises strictly from the assumption that backing assets are perpetuities. If monetary production is backed by finite-maturity assets (e.g., T-period Treasury bonds), the price is a finite geometric sum, well-defined for any $\omega > 0$, completely eliminating this parameter restriction. The saddle-node bifurcation at $\tilde{q}^*$ remains fully operative at high interest rates under finite maturities; the perpetuity specification is used here solely for closed-form elegance.

0.107. When backing assets are fully performing ($\tilde{q} = 1$), the model generates an aggregate equity buffer of $E/A \approx 23.9\%$—the maximum capital cushion the monetary system can generate from perfect-quality assets under this parameterization.

This threshold structure can be illustrated using the balance sheet magnitudes from the 2023 U.S. banking episode analyzed by Jiang et al. (2024). They document that following the 2022 monetary tightening cycle, the average U.S. bank experienced mark-to-market asset losses of roughly 10% due to duration risk. While these were discount-rate losses rather than credit defaults, one can use this magnitude to ask a hypothetical structural question: *If* the representative bank had suffered an exogenous cash-flow impairment of this size ($\tilde{q} \approx 0.90$), where would it fall in the theoretical state space? Because $0.90 > \tilde{q}^{**} \approx 0.868$, such a bank would remain in the safe region, consistent with the absence of a generalized systemic collapse.

By contrast, Silicon Valley Bank (SVB) incurred mark-to-market declines of about 15% (see Jiang et al., 2024). If I apply this 15% loss as a hypothetical shock to fundamental cash flows ($\tilde{q} \approx 0.85$), SVB is situated squarely in the fragile region: $0.761 < 0.85 < 0.868$. In this region, fundamental cash flows alone cannot cover liabilities ($F_t < \mathcal{M}_t$), meaning the bank is fundamentally insolvent without the liquidity premium. The fact that SVB's failure ultimately materialized through a depositor run rather than through deterministic failure from fundamentals alone ($\tilde{q} \leq \tilde{q}^*$) perfectly illustrates the mechanics of Theorem 7(ii) for intermediate states of asset quality.

Table 1 shows how the critical threshold $\tilde{q}^*$ varies with each parameter, holding the others at their baseline values. All evaluated parameter combinations strictly satisfy the convergence restriction of Theorem 5.

Parameter	**Low**	**Baseline**	**High**	$\tilde{q}^*$ **(Low)**	$\tilde{q}^*$ **(High)**	$\mathcal{G}$ **(High)**
i (nominal rate)	0.005	0.01	0.015	0.81	0.71	0.16
v (velocity weight)	0.3	0.5	0.7	0.80	0.72	0.15
κ (cash-flow share)	0.04	0.048	0.08	0.91	0.46	0.06
β (discount factor)	0.955	0.96	0.965	0.87	0.65	0.11

Table 1: Sensitivity of the crisis threshold $\tilde{q}^*$ and fragility index $\mathcal{G} = \omega/(\beta\kappa)$ to parameter variation. Baseline values: $\tilde{q}^* \approx 0.76$ and $\mathcal{G} \approx 0.11$. All parameter combinations satisfy the convergence condition $\omega < 1 - \beta$. Higher nominal rates and velocity weights *lower* $\tilde{q}^*$ and *widen* the fragile region ($\mathcal{G}$ increases), making the monetary system structurally more fragile (Theorem 12). The $\mathcal{G}$ column reports the fragility index evaluated at the High parameter value.

The table shows that $\tilde{q}^*$ ranges from approximately 0.46 to 0.91 across parameterizations satisfying the convergence condition. The sensitivity to κ is especially pronounced: when

backing-asset cash flows are small relative to GDP ($\kappa = 0.04$), the upper threshold $\tilde{q}^{**} = (1-\beta)/(\beta\kappa) \approx 1.04 > 1$, implying that the safe region is unreachable. This implies the economy is *permanently fragile* regardless of asset performance.

3.6.2 Systemic Permanent Fragility: The 1929 Banking Sector

The hypothetical SVB illustration demonstrates idiosyncratic entry into an avoidable fragile region (where $\tilde{q}^{**} < 1$). To evaluate the model against a system-wide regime of actual credit impairment, I construct a second quantitative proof of concept calibrated to the aggregate U.S. banking sector on the eve of the Great Depression. This period provides an ideal empirical laboratory because subsequent failures were driven by observable fundamental asset impairment (loan losses) rather than pure duration risk (Calomiris and Mason, 2003).

I set the universal macroeconomic parameters to standard annual targets: a discount factor of $\beta = 0.96$ (**?**) and a nominal interest rate of $i = 0.04$ representing the yield on safe, liquid alternative assets in 1929 (**?**). This yields $\lambda \approx 0.0385$. To capture the "moneyness" of bank deposits, I set the relative transaction velocity to $v = 0.4$, consistent with empirical estimates of the liquidity premium of near-money assets (**?**). Under these parameters, the effective liquidity intensity is $\omega \approx 0.0154$, generating a liquidity production coefficient of $\Phi \approx 14.86$.

To pin down the fundamental cash-flow capacity κ, I follow the standard structural approach of targeting an observable macroeconomic moment. Calomiris and Mason (2003) report an aggregate Net Worth to Total Assets ratio of 14.9% for Federal Reserve member banks in December 1929. Matching this equity ratio requires calibrating the fundamental capacity to $\kappa = 0.030$. This internally generated value perfectly aligns with historical banking data; the **?** reports that the net return on earning assets for member banks (after operating expenses and charge-offs) averaged approximately 3% in the late 1920s.

This parameterization yields two profound structural insights regarding the boundaries of the state space:

1. **Permanent Fragility ($\tilde{q}^{**} > 1$):** The threshold for absolute safety requires fundamental cash flows to independently cover all liabilities: $\tilde{q}^{**} = (1-\beta)/(\beta\kappa) \approx 1.389$. Because maximum physical quality is strictly bounded at $\tilde{q}_t \leq 1$, the banking system mathematically could not reach the safe region. The 1929 banking sector was in a regime of permanent structural fragility, sustained entirely by the endogenous liquidity premium $\Phi\kappa$.

2. **The Insolvency Threshold ($\tilde{q}^*$):** The deterministic run threshold evaluates to:

$$\tilde{q}^* = \frac{1 - \beta - \omega}{\beta\kappa} \approx 0.854$$

The model predicts that if loan-portfolio quality rots to roughly 85 cents on the dollar, the monetary equilibrium unravels completely. This mathematically derived threshold aligns exceptionally well with historical recovery rates. Anari et al. (2005) document that depositors in banks fully liquidated during the 1930–1933 period recovered an average of 75 to 80 cents on the dollar, confirming the model's structural prediction for fundamental collapse.

Discussion: Two Regimes of Fragility. The contrast between the SVB illustration and the 1929 aggregate calibration reveals a core structural insight regarding the model's state space. In the SVB case, the bank entered the fragile region, but absolute fundamental safety ($\tilde{q}^{**} < 1$) remained mathematically attainable with sufficient capitalization ($\kappa = 0.048$). By contrast, the 1929 calibration yields $\tilde{q}^{**} = 1.389 > 1$. Because fundamental quality cannot exceed unity, the safe region was structurally unreachable for the 1929 banking system. This establishes two distinct failure modes within the same theoretical framework: idiosyncratic entry into an avoidable fragile region, versus a system-wide regime of permanent fragility where survival is strictly contingent on sustaining the endogenous liquidity premium.

4 Unification of Crisis Types

The monetary production technology, as defined in Section 3, is agnostic about the specific identity of the assets underlying the liquid media. This structural agnosticism implies that the identical crisis mechanism operates whenever the fundamental quality $\tilde{q}_t$ of the backing assets deteriorates—regardless of whether those backing resources are held on a private balance sheet, physically embodied in the medium itself, or constituted by sovereign taxing capacities. In this section, I formalize this insight by demonstrating that banking crises, currency crises, sovereign debt crises, and historical debasement panics are all manifestations of the single threshold condition $\tilde{q}_t < \tilde{q}^*$, applied across distinct institutional arrangements for monetary production.

4.1 Unification Framework

I first establish a general result: any monetary arrangement satisfying three sufficient conditions can be represented as a monetary production technology and inherits the crisis

threshold of Theorem 7. The four canonical crisis types then follow as a corollary.

Definition 13 (Monetary Arrangement)**.** A **monetary arrangement** is a quintuple $\mathfrak{A} = (i, \mathcal{J}_B, \mathcal{J}_L, v, \tau)$ where:

(i) i is an agent (the monetary producer);

(ii) $\mathcal{J}_B$ is a nonempty set of objects (the **backing assets**) *underlying* the monetary instruments supplied by i — that is, objects whose quantity, quality, or value determines (i) the fundamental worth of the liquid media, (ii) the supply of liquid media that i can credibly sustain, and (iii) the purchasing power that the liquid media anchor. Ownership of the backing assets by i is **not required**: what matters is the economic relationship between the backing resource and the liquid medium, not the legal title;

(iii) $\mathcal{J}_L$ is a nonempty set of objects (liquid media) supplied by i;

(iv) $v : \mathcal{J}_L \to (0, 1]$ assigns a velocity weight to each liquid medium; and

(v) $\tau \in \{\textsc{claim}, \textsc{embodied}\}$ is the **backing type**, indicating whether the liquid media are claims *on* the backing assets (as with bank deposits or banknotes) or physically *embody* the backing assets (as with commodity coins).

The backing relationship in $\mathcal{J}_B$ takes three canonical forms depending on the arrangement: *balance-sheet backing*, in which i holds the backing assets on its own balance sheet (commercial banking: loans back deposits); *resource backing*, in which the backing assets exist independently of i's balance sheet but constrain the supply of liquid media through a physical or legal link (gold coinage: the aggregate bullion stock underlies coins, constraining their supply and anchoring their purchasing power, even if the mint does not own the bullion at the time of striking); and *institutional backing*, in which the backing "assets" are capacities or powers rather than balance-sheet objects (fiat currency: the government's fiscal capacity, taxing power, and legal-tender status underlie the currency's value without appearing on any single institution's balance sheet). The Unification Proposition (Theorem 14) applies under all three forms.

Proposition 14 (Unification)**.** *Let $\mathfrak{A} = (i, \mathcal{J}_B, \mathcal{J}_L, v, \tau)$ be a monetary arrangement satisfying the following **sufficient** conditions in the CIA economy with log utility. (I do not claim these conditions are necessary; the question of minimal conditions for the three-region structure under general liquidity constraints is left for future work.)*

(A1) Backing identity. *The economic value of the backing assets underlying i's liquid media equals the value of those liquid media plus a residual claim (equity buffer):*

$$\underbrace{A_t}_{\substack{\text{value of backing} \\ \text{assets underlying } \mathcal{J}_L}} = \underbrace{\mathcal{M}_t}_{\text{liquid media}} + E_t$$

where A_t, $\mathcal{M}_t$, and E_t are measured consistently with the backing type τ and the form of backing:

- *Balance-sheet (CLAIM, i holds $\mathcal{J}_B$ on its own books): A_t is the market value of the asset portfolio, $\mathcal{M}_t$ is the outstanding nominal liabilities, and $E_t \geq 0$ is i's equity.*
- *Resource (EMBODIED, $\mathcal{J}_B$ underlies $\mathcal{J}_L$ without balance-sheet ownership): $A_t = p_t^{metal} \cdot S_t$ is the aggregate intrinsic value of the backing stock[5] (S_t = total metal stock available to the mint); $\mathcal{M}_t$ is the aggregate nominal face value of coins outstanding; and $E_t = A_t - \mathcal{M}_t$ is the aggregate seigniorage buffer (which equals zero for fully backed coinage at issuance). The physical constraint is $N_t \cdot s_t \leq S_t$—total metal embodied in coins cannot exceed the available bullion stock.*
- *Institutional (CLAIM, $\mathcal{J}_B$ = fiscal/legal capacity): A_t is the present discounted value of the government's primary surpluses (the "fiscal backing"); $\mathcal{M}_t$ is the monetary base; and $E_t = A_t - \mathcal{M}_t$ is the fiscal buffer. The balance sheet is consolidated across the fiscal authority and the central bank.*

(A2) Liquidity entry. *The liquid media enter households' liquidity constraints with a positive velocity weight:*

$$\frac{\partial L_t^h}{\partial (a_t^{k,h} p_t^k)} = v_t^k > 0 \qquad \forall\, k \in \mathcal{J}_L$$

(A3) Fundamental separability. *Backing-asset quality $\tilde{q}_t \in (0, 1]$ is an exogenous fundamental state variable, independent of whether the liquid media circulate as money. The market value of the backing assets admits the decomposition $A_t = F_t + LP_t$, where the fundamental value F_t is determined by $\tilde{q}_t$ and exogenous fundamentals, and the liquidity premium LP_t is endogenous to the monetary equilibrium. The definition of $\tilde{q}_t$ and the form of F_t depend on the backing type:*

[5]While physical bullion does not pay a literal cash dividend, its intrinsic value can be mapped to my general asset-pricing framework by interpreting y^A as the strictly positive, non-monetary convenience yield (or implicit rental rate) derived from the metal's industrial or artistic utility. The present value of this constant utility flow is $F_t = (\tilde{q}_t \cdot y^A)/r$, where $\tilde{q}_t = s_t/s_0$ captures physical debasement, thereby justifying the application of the perpetuity multiplier Φ to commodity money.

- *CLAIM ($\tau =$ CLAIM): $\tilde{q}_t$ is the ratio of impaired expected cash flows to the unimpaired (promised) benchmark cash flows. The fundamental value $F_t = \tilde{q}_t \cdot y^A/r$ is* ***evaluated independently of whether the liquid media circulate as money****—i.e., using the SDF μ alone, without the liquidity kernel Λ (fundamental separability; see* **??***). The liquidity premium $LP_t = \Phi \cdot \tilde{q}_t \cdot y^A$ is present if and only if the liquid media are accepted as money.*

- *EMBODIED ($\tau =$ EMBODIED): $\tilde{q}_t = s_t/s_0$, where s_t is the current intrinsic physical content (e.g., metal weight) of the liquid medium and s_0 is the full-weight statutory standard. The fundamental (intrinsic) value is $F_t = \tilde{q}_t \cdot p_t^{face}$, and the liquidity premium is $LP_t = (1 - \tilde{q}_t) \cdot p_t^{face} \geq 0$. The liquid medium is fully backed at $\tilde{q}_t = 1$ and debased at $\tilde{q}_t < 1$.*

Then:

(i) $\mathfrak{A}$ defines a monetary production technology $\mathcal{T}_t$ in the sense of Theorem 3.

(ii) The crisis threshold of Theorem 7 applies: there exist $\tilde{q}^* < \tilde{q}^{**}$ generating the three-region equilibrium structure (safe, fragile, crisis), with

$$\tilde{q}^* = \frac{1 - \beta - \omega}{\beta\kappa} \tag{39}$$

where $\omega = \lambda\bar{v}$ is the effective liquidity intensity of the liquid media and κ is the fundamental cash-flow (or intrinsic-value flow) capacity of the backing assets.

(iii) The amplification mechanism of Theorem 9 operates within the fragile region, and the saddle-node bifurcation at $\tilde{q}^*$ is generic.

Proof. (i). By (A2), the liquid media $\mathcal{J}_L$ enter L_t^h with positive weights, so agent i is a monetary producer (Theorem 4). By (A1), there exist backing resources (held on the balance sheet, embodied in the medium, or constituted by institutional capacity) whose aggregate value A_t constrains the stock of liquid media $\mathcal{M}_t$. Ownership of $\mathcal{J}_B$ by i is not required; what is required is the *economic* constraint $\mathcal{M}_t \leq A_t + E_t$. By Theorem 3, the mapping from the backing-resource value A_t to $\ell_t = \max x_t^{1,h}$ subject to $L_t^h \geq 0$ is a monetary production technology under either backing type.

(ii) for $\tau =$ CLAIM (balance-sheet or institutional backing): By (A1), $A_t = \mathcal{M}_t + E_t$. By (A3), $A_t = F_t + LP_t$ where $F_t = \tilde{q}_t \cdot y^A/r$ is determined by exogenous fundamentals, independent of whether $\mathcal{M}_t$ circulates as money (fundamental separability). The liquidity premium $LP_t = \Phi(\omega, \beta) \cdot \tilde{q}_t \cdot y^A$ is present if and only if the liquid media are accepted. The Euler equation

(12) with the CIA constraint yields the steady-state kernel $\Lambda = 1/(1-\omega)$, where $\omega = \lambda\bar{v}$. At the solvency boundary $E_t = 0$: $F_t + LP_t = \mathcal{M}_t$. Substituting $F_t + LP_t = \tilde{q}_t \cdot y^A(1/r + \Phi)$ and money market clearing $\mathcal{M}_t/P_t = Y_t$ then yield $\tilde{q}^* = (1-\beta-\omega)/(\beta\kappa)$.

(ii) for $\tau =$ EMBODIED (resource backing): The backing resource is the aggregate bullion stock S_t. The mint need not own S_t; the constraint is $N_t \cdot s_t \leq S_t$ (total metal embodied in coins cannot exceed the available stock). The mapping is: $S_t \to N_t$ coins $\to \ell_t$ liquidity services. Each coin's market value decomposes as:

$$p_t^{coin} \;=\; \underbrace{p_t^{metal} \cdot s_t}_{F_t = \tilde{q}_t \cdot p_t^{face}} \;+\; \underbrace{p_t^{face} - p_t^{metal} \cdot s_t}_{LP_t \geq 0}$$

where the substitution $p_t^{metal} \cdot s_t = \tilde{q}_t p_t^{face}$ follows directly from the definition $\tilde{q}_t = s_t/s_0$ under the standard assumption that a full-weight coin's melt value equals its nominal face value ($p_t^{metal} \cdot s_0 = p_t^{face}$). The liquidity premium $LP_t \geq 0$ is the acceptability premium above melt value. At full backing ($s_t = s_0$, no debasement), $\tilde{q}_t = 1$ and LP_t is the pure seigniorage of the monetary system. The equilibrium constraint requires that coins be accepted at par: $p_t^{coin} = p_t^{face}$, which implies $LP_t = p_t^{face}(1 - \tilde{q}_t) \geq 0$. The crisis threshold is reached when $\tilde{q}_t < \tilde{q}^*$: the intrinsic content is so far below face value that no positive liquidity premium can sustain par acceptance, and Gresham's Law operates.

The three-region structure persists:

- Safe ($\tilde{q}_t > \tilde{q}^{**}$): intrinsic value alone sustains acceptance; the liquidity premium is positive but not load-bearing.
- Fragile ($\tilde{q}^* < \tilde{q}_t < \tilde{q}^{**}$): the coin circulates at par in the monetary equilibrium (the liquidity premium sustains it) but not in the autarkic equilibrium.
- Crisis ($\tilde{q}_t < \tilde{q}^*$): even the maximum liquidity premium cannot sustain par acceptance; Gresham's Law drives debased coins out of circulation.

(iii). The amplification mechanism and saddle-node bifurcation hold under all three backing forms by the same argument as Theorem 7(iii): the self-referential fixed point (acceptance $\Rightarrow$ liquidity premium $\Rightarrow$ backing-resource value $\Rightarrow$ acceptance) exists for any arrangement satisfying (A1)–(A3), and the boundary at $\tilde{q}^*$ is a generic collision of the monetary and autarkic equilibria. The ownership structure of $\mathcal{J}_B$ is irrelevant to this mechanism: what matters is that A_t constrains $\mathcal{M}_t$, whatever the legal form of that constraint. □ □

Remark 15 (Contagion and Endogenous Backing). The fundamental-separability requirement (A3) intentionally excludes cases where the backing assets are themselves monetary instruments

(e.g., interbank deposits or repo collateral). In such arrangements, the fundamental value of the backing asset depends intrinsically on whether the underlying instrument circulates as money, violating the strict exogeneity of $\tilde{q}_t$. This boundary condition formally defines the transmission mechanism for financial contagion: monetary system failures propagate when one producer's liquid liabilities serve as another producer's backing assets. I formally relax (A3) to analyze these dynamics in Section 4.7 (Twin Crises).

With the general result in hand, I apply it to the three canonical crisis types.

Corollary 16 (Classification of Financial Crises)**.** *The following monetary arrangements satisfy conditions (A1)–(A3) of Theorem 14 and therefore inherit the three-region equilibrium structure with crisis threshold $\tilde{q}^*$:*

Crisis type	***Type*** τ	***Backing form***	***Backing*** $(\mathcal{J}_B)$	***Liquid media*** $(\mathcal{J}_L)$	***Quality*** $\tilde{q}_t$
Banking crisis	CLAIM	*Balance-sheet*	*Loan portfolio*	*Deposits, notes*	$\tilde{q}_t^b = F_t^b/F_t^{\text{n}}$
Currency crisis	CLAIM	*Balance-sheet*	*FX reserves, gold*	*Domestic currency*	$\tilde{q}_t^R = F_t^R/F$
Sovereign debt	CLAIM	*Institutional*	*Fiscal capacity*	*Gov't-backed money*	$\tilde{q}_t^g = F_t^g/F_t^{\text{l}}$
Debasement crisis	EMBODIED	*Resource*	*Bullion stock* S_t	*Coin at face value*	$\tilde{q}_t^{deb} = s_t/s_0$

The "Backing form" column indicates the ownership relationship: balance-sheet (the producer holds $\mathcal{J}_B$ on its own books), institutional (the backing is a legal/fiscal capacity rather than a held asset), and resource (the backing stock underlies the liquid media through a physical constraint, without requiring ownership by the producer).
In each case, the crisis occurs when fundamental quality $\tilde{q}_t$ falls below $\tilde{q}^* = (1-\beta-\omega)/(\beta\kappa)$.

Proof. I verify (A1)–(A3) for each arrangement.

Banking crisis ($\tau = \text{CLAIM}$). (A1): $A_t^b = D_t^b + E_t^b$. (A2): Deposits enter L_t^h with $v_t^{D^b} > 0$. (A3): Loan fundamental value F_t^b (PDV of expected repayments) is determined by borrower capacity, strictly independent of whether deposits circulate. ✓

Currency crisis ($\tau = \text{CLAIM}$). (A1): $A_t^R = N_t + E_t^{CB}$. (A2): Domestic currency enters L_t^h with $v_t^N > 0$. (A3): Reserve fundamental value F_t^R is determined by the anchor country's monetary policy and exchange fundamentals, independent of whether domestic currency is accepted. ✓

Sovereign debt crisis ($\tau = \text{CLAIM}$). (A1): $A_t^g = \mathcal{M}_t^g + E_t^g$. (A2): Government-backed money enters L_t^h with $v_t^g > 0$. (A3): F_t^g (PDV of primary surpluses) is determined by fiscal

fundamentals, independent of whether government money circulates.[6] ✓

Debasement crisis ($\tau =$ EMBODIED, resource backing). (A1): The backing resource is the aggregate bullion stock S_t; the mint need not own it. The resource constraint is $N_t s_t \leq S_t$. $A_t = p_t^{metal} S_t$ (total intrinsic value of available metal); $\mathcal{M}_t = p_t^{face} N_t$ (face value of coins outstanding); $E_t = A_t - \mathcal{M}_t \geq 0$ (seigniorage buffer, equals zero at full backing $s_t = s_0$). (A2): Coins enter L_t^h with $v_t^{coin} > 0$ so long as they are accepted at par; the velocity weight collapses to zero in the crisis region. (A3): $\tilde{q}_t = s_t/s_0$, the physical ratio of current intrinsic content to the full-weight standard. Debasement is the mint's choice to reduce s_t while holding p_t^{face} fixed, directly reducing $\tilde{q}_t$. The crisis threshold $\tilde{q}^*$ is reached when the liquidity premium cannot sustain par acceptance, triggering Gresham's Law. ✓

By Theorem 14, each arrangement inherits the full three-region structure under its respective (A3) case. □ □

4.2 Banking Crises

When depositors in 1930s America lined up outside their banks, they were not irrational—they had correctly deduced that the loans backing their deposits were fundamentally impaired (Calomiris and Mason, 2003).

In the banking-crisis specification, the backing assets are the bank's loan portfolio. Quality $\tilde{q}_t^b$ measures the fraction of promised repayments that borrowers will actually deliver $(F_t^b / F_t^{\max,b})$. In normal times, $\tilde{q}_t^b$ is high: loans are performing, and the bank's equity cushion absorbs losses. The liquidity premium on the loan portfolio is positive because deposits are accepted as money, and this premium reinforces bank solvency.

The crisis chain operates as follows:

1. An adverse real shock (recession, asset price decline, sectoral collapse) reduces borrower repayment capacity, lowering $\tilde{q}_t^b$.

2. As $\tilde{q}_t^b$ declines, bank equity $E_t^b = A_t^b - D_t^b$ erodes. The bank enters the fragile region $(\tilde{q}^* < \tilde{q}_t^b < \tilde{q}^{**})$, where solvency depends entirely on the liquidity premium.

3. If $\tilde{q}_t^b$ crosses $\tilde{q}^*$, even the maximum liquidity premium cannot restore solvency. The economy coordinates on the non-monetary equilibrium, in which deposits cease to serve

[6] Condition (A3) for sovereign debt requires that seigniorage revenue is negligible relative to primary surpluses—equivalently, that fiscal backing is determined by taxation capacity rather than monetary policy. Under the Fiscal Theory of the Price Level (Cochrane, 2023), seigniorage enters the fiscal backing equation and (A3) fails. In that case, the sovereign crisis is better mapped to the EMBODIED case ($\tau =$ EMBODIED) with "intrinsic value" reinterpreted as the fiscal backing ratio—the connection noted in Theorem 17. For advanced economies where seigniorage is small relative to tax receipts, (A3) remains a robust approximation.

as transaction media ($v_t^{D^b} = 0$): no household accepts them because no household expects others to accept them (see Theorem 10). The liquidity premium collapses, and inside money production ceases.

4. The total money supply contracts from $M_t = M^{out} + D_t^b$ to $M_t = M^{out}$. The shadow cost of liquidity λ_t spikes.

A crucial distinction from Diamond and Dybvig (1983): the banking crisis in my model does not require a bank run. The collapse occurs because fundamentals deteriorate past $\tilde{q}^*$, not because depositors fail to coordinate. This is consistent with Baron et al. (2021), who show that many systemic banking crises—including the early stages of the 2007–2008 crisis and the 2010–2012 European crisis—featured gradual asset deterioration rather than sudden depositor panics. In my framework, runs *can* occur in the fragile region (as sunspot equilibria), but the crisis region is reached through fundamentals alone.

The three-region prediction—that fundamental insolvency and run-driven fragility are distinct mechanisms with different observable signatures—receives strong empirical support from recent work on U.S. banking history. Correia et al. (2026a) study 160 years of bank failures (1865–2023) and decompose each into solvency and liquidity components. Their central finding is that approximately three-quarters of all bank failures are driven primarily by fundamental asset deterioration (my crisis region, $\tilde{q}_t < \tilde{q}^*$), with depositor runs playing an accelerating rather than causal role. Creditors of fundamentally insolvent banks recovered approximately 72 cents on the dollar, while creditors of illiquid-but-solvent failures recovered roughly 92 cents—a 20-percentage-point gap that corresponds to the economic distinction between the crisis and fragile regions. In a companion study, Correia et al. (2026b) examine bank runs that do *not* lead to failure and find that runs frequently strike fundamentally strong institutions, which survive precisely because $\tilde{q}_t > \tilde{q}^{**}$ places them in the safe region where solvency does not depend on the liquidity premium. Taken together, these findings confirm the model's core prediction: the run is neither necessary nor sufficient for crisis; what determines survival is whether backing-asset quality lies above or below the threshold.

4.3 Currency Crises

When the Bank of England abandoned the gold standard in September 1931, it was acknowledging that its gold reserves could no longer fundamentally sustain the liquid liabilities denominated in sterling (Eichengreen, 1992).

In the currency-crisis specification, the monetary producer is a central bank operating a fixed exchange rate or metallic standard. The backing assets *underlying* the domestic currency are foreign exchange reserves (or gold); ownership by the central bank is required in

this case (balance-sheet CLAIM), since the peg commitment is precisely a promise to convert domestic currency into the reserve asset on demand.

The CIA constraint for the domestic economy takes the form:

$$a_t^{M^*,h} p_t^{M^*} e_t + a_t^{N,h} p_t^N v_t^N \geq x_t^{1,h} \tag{40}$$

where M^* is foreign currency, e_t is the exchange rate, and N is domestic currency. The central bank's balance sheet is:

$$A_t^R = N_t + E_t^{CB} \tag{41}$$

where $A_t^R = R_t \cdot e_t$ is the domestic value of foreign exchange reserves. Quality is $\tilde{q}_t^R = F_t^R / F_t^{\max,R}$, where F_t^R reflects the fundamental, sustainable reserve capacity generated by current account dynamics, relative to a fully backed benchmark $F_t^{\max,R}$.

The crisis threshold is:

$$\tilde{q}^{*,FX} = \frac{1 - \beta - \omega}{\beta \kappa^R} \tag{42}$$

where $\omega = \lambda v^N$ is the effective liquidity intensity of the domestic currency, and κ^R is the cash-flow return on reserves (e.g., interest on foreign-currency bonds) as a share of output. When fundamental reserves are depleted ($\tilde{q}_t^R < \tilde{q}^{*,FX}$), the central bank cannot credibly maintain the peg, and the domestic currency collapses.

The classical speculative attack models (Krugman, 1979; Obstfeld, 1996) map naturally onto my framework. First-generation attacks (Krugman, 1979) correspond to fundamental deterioration into the crisis region: unsustainable fiscal deficits drain reserves until $\tilde{q}_t^R < \tilde{q}^{*,FX}$, making collapse inevitable. Second-generation attacks (Obstfeld, 1996) correspond to the fragile region: reserves are sufficient in the "good" equilibrium ($\tilde{q}_t^R > \tilde{q}^{*,FX}$) but not if agents coordinate on the "bad" equilibrium (withdrawing reserves and destroying the liquidity premium). My framework unifies both generations by providing the exact boundaries ($\tilde{q}^*$ and $\tilde{q}^{**}$) between the three regions.

The modern currency crisis (reserve depletion under a pegged rate) and the ancient currency crisis (coinage debasement) are both captured by Theorem 16, but under *different backing types*. Reserve crises are $\tau = \text{CLAIM}$: the currency is a claim on reserves, and the crisis occurs when the claim cannot be honored. Debasement crises are $\tau = \text{EMBODIED}$: the coin *is* the metal, and the crisis occurs when the physical metal content falls below the acceptance threshold. The shared mechanism is the exact same collapse of the self-referential fixed point at $\tilde{q}_t = \tilde{q}^*$; the institutional form differs. I develop the debasement case formally in Section 4.5.

4.4 Sovereign Debt Crises

When Greek government bond yields spiked above 30% in 2012, the euros deposited in Greek banks became suspect—because those banks held the very government bonds that the market was repricing (Farhi and Tirole, 2018).

The sovereign debt crisis is a hybrid: it combines elements of both banking and currency crises. The backing assets are government bonds, which serve dual roles. First, they back government-issued or government-guaranteed money (directly, as with Treasury bills used as collateral, or indirectly, through the fiscal backing of the central bank). Second, they constitute a large fraction of banking-system assets, so their deterioration simultaneously triggers a banking crisis.

In the tractable specification, let the banking system hold government bonds B^g as backing assets:

$$A_t^{bank} = \alpha_g B_t^g + (1 - \alpha_g) L_t \tag{43}$$

where α_g is the share of government bonds in bank assets and L_t is private loans. The fundamental quality of the government bond portfolio is evaluated strictly on expected versus promised cash flows:

$$\tilde{q}_t^g = \frac{\text{PDV of expected primary surpluses}}{\text{PDV of promised bond cash flows}} \tag{44}$$

Fiscal deterioration (declining primary surpluses, rising debt issuance) reduces $\tilde{q}_t^g$ exogenously. The effective banking-system quality is:

$$\tilde{q}_t^{eff} = \alpha_g \tilde{q}_t^g + (1 - \alpha_g) \tilde{q}_t^L \tag{45}$$

A sovereign debt crisis occurs when $\tilde{q}_t^{eff} < \tilde{q}^*$. The transmission is amplified by the sovereign-bank nexus: fundamental fiscal deterioration reduces $\tilde{q}_t^g$, which pushes the banking system into the fragile region, triggering deposit flight, the collapse of the liquidity premium, and ultimately forcing banks into insolvency.

4.5 Debasement Crises

The Roman denarius contained approximately 4.5 grams of silver under Augustus. By the reign of Gallienus (260–268 AD), it contained roughly 0.02 grams—a debasement to less than half a percent of the original (Harl, 1996). Long before the final collapse, Roman merchants quoted prices in gold, accepted only full-weight coins, or demanded large premiums on debased ones. This is Gresham's Law in operation: bad money drives out good. In my framework, it is the crisis region of the EMBODIED monetary arrangement.

The debasement arrangement. The monetary producer is the mint (controlled by the sovereign). The backing assets underlying the coinage system are the aggregate bullion stock S_t—the total supply of silver or gold available in the economy. The mint need not own this stock: the economic relationship is that the bullion stock constrains coin supply through the physical accounting identity $N_t \cdot s_t \leq S_t$ (coins struck times metal per coin cannot exceed available metal). This is resource backing: the bullion stock *underlies* the coins by determining their intrinsic value, constraining how many can credibly circulate, and anchoring their purchasing power relative to non-monetary uses of metal.

The liquid medium is the coin at its official face value p^{face}—a legal fiction enforced by the sovereign through legal-tender laws and tax-receivability, which holds rigid while the intrinsic metal value $p_t^{metal} \cdot s_t$ fluctuates or is actively debased. This rigidity is exactly what forces $\tilde{q}_t$ downward toward the crisis threshold. To satisfy fundamental separability (A3), quality per coin is defined strictly by physical content:

$$\tilde{q}_t^{deb} = \frac{s_t}{s_0} \tag{46}$$

where s_t is current metal weight and s_0 is the full-weight statutory standard. Under full backing ($\tilde{q}_t^{deb} = 1$), the coin's metal content equals its face value and no seigniorage is extracted. Debasement reduces s_t while holding p_t^{face} fixed, lowering $\tilde{q}_t^{deb}$. Crucially, the mapping from backing resource to liquid medium is: $S_t \to N_t$ coins $\to \ell_t$ liquidity services per period—the same production structure as Theorem 3, with the bullion stock playing the role of the backing-asset portfolio.

The three-region structure. Applying Theorem 14 with $\tau =$ EMBODIED:

1. **Safe** ($\tilde{q}_t^{deb} > \tilde{q}^{**}$): Coins circulate at par and are accepted without inspection. The liquidity premium of the coin (its acceptability above melt value) is robust even if the metal premium fluctuates.

2. **Fragile** ($\tilde{q}^{*} < \tilde{q}_t^{deb} < \tilde{q}^{**}$): Coins still circulate but merchants begin inspecting weights, charging discounts, or quoting prices in metal rather than face value. The monetary equilibrium persists as long as enough agents accept the coin at face value.

3. **Crisis** ($\tilde{q}_t^{deb} < \tilde{q}^{*}$): Full-weight coins are hoarded or melted. Only debased coins circulate (Gresham), and these are rejected at par. The effective money supply collapses to whatever full-weight coins remain in circulation.

Unlike the other three crisis types—where quality deterioration is an exogenous shock to backing assets—debasement is an endogenous choice by the monetary producer. The sovereign

deliberately reduces s_t to extract seigniorage revenue $\pi_t^{mint} = (p_t^{face} - p_t^{metal} s_t) \cdot \Delta coins_t$. This has an important implication for the crisis dynamics: the producer has foreknowledge of the quality path and can anticipate the threshold crossing. Rational producers will debase up to $\tilde{q}^*$ and stop—but time-inconsistency problems (the temptation to extract one more period of seigniorage) can push $\tilde{q}_t^{deb}$ below the threshold, triggering the collapse they sought to avoid. Sargent and Velde (2002) document this mechanism in the medieval European experience with small change.

The debasement case is developed here as an application of the unification framework (conditions (A1)–(A3) with $\tau = \text{EMBODIED}$) using the same CIA/log-utility tractable specification. A complete formal model of commodity coinage—with an explicit minting technology, bullion market equilibrium, endogenous seigniorage choice, and the Gresham adverse-selection mechanism—is beyond the scope of this paper and is left for future work. What the framework establishes is that the three-region equilibrium structure and the crisis threshold $\tilde{q}^*$ apply to commodity coinage under the same structural conditions as for claim-based monetary arrangements.

Remark 17 (Connection to modern fiat debasement). The EMBODIED backing type also applies, by analogy, to fiat currencies under discretionary monetary policy. When a central bank creates money to finance fiscal deficits, it is reducing the "intrinsic" fiscal backing of the currency—the PDV of expected tax revenues relative to the monetary base. In this reading, $\tilde{q}_t^{deb}$ measures the fiscal backing ratio (a concept central to the Fiscal Theory of the Price Level). The debasement crisis in fiat systems is hyperinflation: the currency is accepted at par until the fiscal backing collapses below $\tilde{q}^*$, at which point velocity explodes and the currency is abandoned for alternatives (barter, foreign currency, or commodity money).

4.6 Inside versus Outside Money

Following Lagos (2006), outside money is defined as money that is either fiat or backed by an asset not in zero net supply, and therefore constitutes a net asset for the private sector. By contrast, inside money is a privately issued liability backed by private credit that circulates as a medium of exchange and is in zero net supply within the private sector. The conventional distinction thus centers on aggregate net supply and balance-sheet accounting.

The unification proposition, however, shifts attention from net supply to production structure. It shows that both inside and outside money are generated through the same underlying mathematical mechanism: in each case, an asset with sufficiently high fundamental value is transformed into circulating media subject to a solvency constraint and a participation threshold. Institutional differences—whether the backing asset is publicly or privately issued,

whether it is in zero net supply, and which agents internalize the relevant constraint—affect the identity of inputs and the equilibrium region in which the economy operates, but not the core production mapping itself. In this sense, inside and outside money are alternative institutional implementations of a single monetary production technology, differentiated by backing structure and threshold parameters rather than by distinct analytical foundations.

Proposition 18 (Inside vs. Outside Money Production)**.** *Both inside money (e.g., bank deposits backed by private assets on the issuer's balance sheet) and outside money (e.g., fiat currency with institutional backing, or commodity money with resource backing) utilize the identical monetary production technology $\mathcal{T}_t$. They differ only in the set of assets* underlying *the liquid media ($\mathcal{J}_B$) and the form of that backing:*

(i) ***Inside money:*** *$\mathcal{J}_B = \{$private loans, securities$\}$ (balance-sheet CLAIM backing). Quality $\tilde{q}_t^{in}$ is determined exogenously by private-sector solvency (the ratio of expected to promised cash flows).*

(ii) ***Outside money (fiat):*** *$\mathcal{J}_B = \{$fiscal capacity, taxing power$\}$ (institutional CLAIM backing). Quality $\tilde{q}_t^{out}$ is determined by the fundamental fiscal backing ratio (the present discounted value of expected primary surpluses relative to a fully backed benchmark).*

(iii) ***Outside money (commodity):*** *$\mathcal{J}_B = \{$bullion stock$\}$ (resource EMBODIED backing). Quality $\tilde{q}_t^{out}$ is determined strictly by the physical ratio of current intrinsic content to the full-weight standard (s_t/s_0).*

(iv) ***Asymmetric Fragility:*** *The crisis thresholds $\tilde{q}^{*,in}$ and $\tilde{q}^{*,out}$ generally differ because the fundamental cash-flow capacities of the backing assets differ ($\kappa^{in} \neq \kappa^{out}$). The two sectors can partially substitute: if inside money collapses ($\tilde{q}_t^{in} < \tilde{q}^{*,in}$) but outside money remains fundamentally sound ($\tilde{q}_t^{out} > \tilde{q}^{*,out}$), the economy contracts but does not collapse into complete autarky.*

Proof. Parts (i)–(iii) follow directly from Theorem 3 and Theorem 14 applied to each money type. The backing form determines which branch of the (A1) identity applies: *balance-sheet* backing for inside money, *institutional* backing for fiat outside money, and *resource* backing for commodity outside money.

For (iv), the partial substitutability follows directly from the CIA constraint. Let the aggregate money supply be $\mathcal{M}_t = M_t^{out} + D_t^{in}$. When inside money collapses ($D_t^{in} \to 0$ due to $\tilde{q}_t^{in} < \tilde{q}^{*,in}$), the economy is forced to operate on the outside money stock M_t^{out} alone. The shadow cost of liquidity λ_t spikes, and aggregate output falls. The aggregate welfare loss is the difference between the fully monetary equilibrium and the outside-money-only

equilibrium, the severity of which depends entirely on the fundamental capacity of the outside money to anchor the transaction need. □ □

4.7 Twin Crises

Reinhart and Rogoff (2009) document that banking and currency crises frequently co-occur—"twin crises" that are deeper and longer than either type alone. Kaminsky and Reinhart (1999) demonstrate that banking crises typically precede currency crises by 12–18 months. The unification framework provides a structural explanation for both the co-occurrence and this specific sequencing.

Corollary 19 (Twin Crises)**.** *If the inside-money and outside-money production technologies share interlinked liquidity constraints or correlated backing-asset quality, then:*

(i) ***Banking-to-currency contagion.*** *A banking crisis ($\tilde{q}_t^{in} < \tilde{q}^{*,in}$) directly increases the probability of a subsequent currency crisis. The collapse of inside money raises the aggregate shadow cost of liquidity λ_t, which mechanically widens the fragile region for outside money: by Theorem 12(i), the crisis threshold $\tilde{q}^{*,out}$ falls while $\tilde{q}^{**,out}$ is unchanged.*

(ii) ***Currency-to-banking contagion.*** *A currency crisis ($\tilde{q}_t^{out} < \tilde{q}^{*,out}$) transmits to the banking system if banks or their borrowers hold unhedged foreign-currency liabilities. Devaluation inflates the domestic-currency value of foreign-denominated debt while leaving domestic-currency revenues unchanged, impairing the repayment ratio and shifting $\tilde{q}_t^{in}$ downward.*

(iii) ***Superadditivity.*** *The twin crisis is deeper than either single crisis because both production technologies fail simultaneously. The economy loses both inside and outside money production, and the resulting welfare loss is superadditive.*

Proof. Part (i): Let aggregate liquidity be $\mathcal{M}_t = M_t^{out} + D_t^{in}$. When inside money collapses ($D_t^{in} \to 0$), the money supply contracts. For a given transaction need, the CIA constraint requires the shadow cost of liquidity λ_t to spike. The effective liquidity intensity of the surviving outside money is $\omega^{out} = \lambda_t v^{out}$. Because ω^{out} increases, the crisis threshold $\tilde{q}^{*,out} = (1-\beta-\omega^{out})/(\beta\kappa^{out})$ strictly *falls* (by Theorem 12(i)), while $\tilde{q}^{**,out} = (1-\beta)/(\beta\kappa^{out})$ is unchanged—widening the fragile region $\mathcal{G}^{out} = \omega^{out}/(\beta\kappa^{out})$ over which a sunspot-driven currency crisis is possible. Simultaneously, if the banking collapse triggers capital flight, the resulting reserve outflows reduce $\tilde{q}_t^{out}$ itself. The outside-money system becomes more fragile

precisely because the inside-money system has failed: the fragile region is wider and the underlying reserve quality is deteriorating.

Part (ii): If the monetary producer abandons the peg (devaluation), the domestic-currency value of foreign-denominated debt spikes while domestic-currency cash flows are unchanged. This acts as a severe exogenous shock to private-sector solvency, reducing the expected repayment ratio $\tilde{q}_t^{in}$. If $\tilde{q}_t^{in}$ falls below $\tilde{q}^{*,in}$, the banking system collapses.

Part (iii): Superadditivity follows from the convexity of the welfare loss in the liquidity shortage. Because $\lambda_t = 1/x_t^{1,h}$ under log utility, the shadow cost is strictly convex in the money supply gap. The welfare loss from both sectors collapsing simultaneously therefore strictly exceeds the sum of the individual losses from isolated failures. □ □

The German banking and currency crisis of 1931 provides a striking historical illustration. Following the Creditanstalt failure in May, German banks faced a massive run on deposits (inside-money crisis: $\tilde{q}_t^{in} < \tilde{q}^{*,in}$). The resulting contraction of domestic liquidity forced the Reichsbank to defend the mark's gold parity by supplying reserves, draining its gold stocks. By July, gold reserves had fallen below the threshold needed to sustain the mark (outside-money crisis: $\tilde{q}_t^{out} < \tilde{q}^{*,out}$). Both production technologies failed simultaneously, requiring capital controls and a banking holiday. The twin crisis was "the most devastating financial crisis of the twentieth century until 2008" (Schnabel, 2004), precisely because it shut down both channels of monetary production.

4.8 The "This Time Is Different" Syndrome

The unification proposition has an important implication for the empirical regularity documented by Reinhart and Rogoff (2009): financial crises recur with remarkable consistency across centuries and institutional settings, despite participants' conviction that "this time is different."

In my framework, each era's monetary innovators believe they have discovered a new class of backing assets ($\mathcal{J}_B$) that is immune to fundamental quality deterioration. John Law's Mississippi Company shares in 1720, the confiscated estates backing French assignats in 1789, the railroad bonds backing trust company deposits in 1873, and the mortgage-backed securities funding the shadow banking system in 2008—each was hailed as a fundamentally safe asset that could sustain unprecedented levels of monetary production. In every case, when fundamental quality $\tilde{q}_t$ inevitably declined, it eventually breached $\tilde{q}^*$, triggering a collapse of the liquidity premium and the sudden cessation of monetary production.

The structural insight is that the crisis boundary $\tilde{q}^*$ is governed by slow-moving economic primitives: the discount factor β (reflecting human time preference), the effective velocity

weight $\bar{v}$ (reflecting transaction technology), and the fundamental cash-flow capacity κ (reflecting real economic returns). None of these parameters shifts dramatically across centuries. What changes is merely the *identity* of the backing assets and the *institutional clothing* of the monetary arrangement. The threshold equation itself ($\tilde{q}^* = (1 - \beta - \omega)/(\beta\kappa)$) remains an invariant law of economic gravity.

The "this time is different" syndrome can therefore be formalized as a persistent cognitive error: the false belief that the $\tilde{q}^*$ constraint has been permanently bypassed—that financial innovation has severed the link between safe liquid media and the fundamental quality of the underlying assets. The model demonstrates why this belief is structurally flawed each time: the crisis threshold is anchored by macroeconomic primitives that financial engineering can obscure, but never erase.

5 Policy Implications

The monetary production framework yields several policy implications. I develop one formally and state the others as remarks, leaving detailed analysis to companion work.

5.1 Monitoring $\tilde{q}_t$, Not Just Leverage

Standard prudential regulation historically focuses on bank leverage ratios and capital adequacy (Acharya et al., 2009). My framework suggests that the more critical variable is backing-asset quality $\tilde{q}_t$—the strictly exogenous measure of fundamental cash-flow capacity, stripped of any endogenous liquidity premium.

Proposition 20 (Regulatory Implication)**.** *A monetary producer can satisfy all standard leverage requirements (e.g., E_t/A_t remains above the regulatory minimum) and nevertheless be in the fragile region ($\tilde{q}^* < \tilde{q}_t < \tilde{q}^{**}$). This occurs because the endogenous liquidity premium inflates the market value of assets (A_t) above their fundamental value (F_t), making the leverage ratio appear healthier than the underlying cash flows can support.*

Proof. The leverage ratio computed at market values is $E_t/A_t = 1 - \mathcal{M}_t/A_t$. By (A3), $A_t = F_t + LP_t$, and since LP_t is positive in a monetary equilibrium, $A_t > F_t$. The market-value leverage ratio is therefore $1 - \mathcal{M}_t/(F_t + LP_t)$, which strictly exceeds the ratio computed at fundamental values, $1 - \mathcal{M}_t/F_t$. The distance between these two metrics is largest precisely when the economy is in the fragile region, where the liquidity premium is structurally load-bearing for solvency. As $\tilde{q}_t$ approaches $\tilde{q}^*$ from above, the market-value leverage ratio remains positive while the fundamental-value measure approaches zero, revealing the extent to which

solvency depends on the endogenous liquidity premium rather than on fundamental cash-flow capacity. □ □

This mathematical observation provides a formal reinterpretation of balance-sheet strength in the run-up to the 2007–2008 crisis. Financial institutions easily satisfied prevailing regulatory capital requirements when measured at market prices, yet those prices incorporated substantial, flighty liquidity premia on mortgage-backed securities (Gorton, 2010). When market liquidity deteriorated, LP_t contracted sharply, revealing that the fundamental value F_t was insufficient to cover outstanding monetary liabilities. Within the present framework, this episode corresponds to a discrete downward movement in fundamental quality $\tilde{q}_t$ across $\tilde{q}^*$, rather than to an abrupt exogenous deterioration in reported leverage per se. The model thus rationalizes how regulatory metrics based on market values systematically fail to detect proximity to the fragile or crisis regions.

A direct implication is that regulators should complement leverage requirements with systematic monitoring of backing-asset quality $\tilde{q}_t$. Operationally, this entails constructing supervisory estimates of fundamental values F_t and comparing them to market valuations A_t, thereby tracking the evolution of the liquidity premium $LP_t = A_t - F_t$. While fundamental valuation is inherently model-based, even coarse supervisory benchmarks can reveal whether reported capital ratios rely primarily on fragile liquidity premia rather than on resilient cash-flow capacity.

The framework also clarifies the role of central bank reserves. In contrast to textbook money-multiplier mechanisms, reserves here do not operate as a mechanical constraint on deposit creation. Instead, they act as a pristine store of fundamental quality for the private banking system. Provided the outside-money system is itself fundamentally sound (i.e., the sovereign avoids the currency and debt crises modeled in Section 4 and remains in the safe region, $\tilde{q}_t^{out} \geq \tilde{q}^{**,out}$), central bank reserves are default-free and maintain a perfect fundamental cash-flow quality of $\tilde{q} = 1$ from the perspective of commercial banks.

Importantly, because reserves provide ultimate transaction services ($v = 1$), they also carry a maximal liquidity premium ($LP_t > 0$), meaning their market value strictly exceeds their fundamental discounted cash flows. Suppose a bank holds a portfolio of risky private assets with quality $\tilde{q}^{\text{risk}} < 1$ alongside reserves with $\tilde{q} = 1$. The bank's aggregate portfolio quality $\tilde{q}_t$ is a weighted average of these components. An increase in reserve holdings raises the bank's aggregate $\tilde{q}_t$ directly, diluting fundamental risk and shifting the institution away from the fragile region, even if total leverage remains unchanged. Reserves therefore stabilize the system by enhancing average fundamental asset quality rather than by acting as quantitative lending constraints.

This perspective provides a clean microfoundation for post-2008 liquidity regulation,

including the Basel III Liquidity Coverage Ratio (LCR). By requiring banks to hold a minimum share of High-Quality Liquid Assets (HQLA) with $\tilde{q} \approx 1$, regulation effectively imposes a lower bound on aggregate portfolio quality $\tilde{q}_t$. In the model, such a constraint ensures that $\tilde{q}_t$ remains permanently above $\tilde{q}^{**}$, confining the economy to the unique, safe monetary equilibrium and eliminating the multiplicity associated with sunspot-driven runs. Liquidity regulation thus operates through fundamental balance-sheet composition—by altering aggregate asset quality—rather than solely through leverage reduction.

5.2 Further Policy Remarks

Three additional implications follow directly from the model but require further formalization in companion work.

Remark 21 (Lender of last resort as equilibrium selection). The LOLR function of Bagehot (1873) acquires a precise mathematical interpretation: it acts as an equilibrium selection mechanism for monetary production in the fragile region. When $\tilde{q}^* < \tilde{q}_t < \tilde{q}^{**}$, the LOLR sustains the monetary equilibrium by signaling that it will absorb any redemption demand, effectively preventing a sunspot-driven shift to the autarkic equilibrium. Crucially, the LOLR cannot rescue the economy from the crisis region ($\tilde{q}_t < \tilde{q}^*$): if backing assets are fundamentally impaired, no amount of liquidity support can restore the monetary equilibrium. This formalizes Bagehot's dictum: lend freely at a penalty rate against *good* collateral ($\tilde{q} > \tilde{q}^*$), but not against fundamentally impaired collateral ($\tilde{q} < \tilde{q}^*$).

Remark 22 (Free banking and decentralized production). Under free banking, monetary production is decentralized: multiple banks issue competing liquid media, each characterized by its own backing-asset quality $\tilde{q}_t^b$ and crisis threshold $\tilde{q}^*$. Decentralization introduces a structural trade-off. On the one hand, idiosyncratic shocks to the backing assets of a single producer need not destabilize the system as a whole, since surviving banks continue to supply liquid media. On the other hand, correlated asset exposures or dense interbank linkages may transmit declines in $\tilde{q}_t^b$ across institutions, pushing several producers below $\tilde{q}^*$ simultaneously and generating systemic contraction.

In a gold-standard free banking system (White, 1984; Selgin, 1988), there is no mechanical money multiplier. Each bank issues private notes backed by a portfolio composed of private loans with $\tilde{q} < 1$ and gold reserves with $\tilde{q} = 1$. The bank's aggregate portfolio quality $\tilde{q}_t^b$ is a weighted average of these components. Gold reserves therefore function as a quality buffer: a sufficiently large reserve share keeps $\tilde{q}_t^b$ above the upper threshold $\tilde{q}^{**}$, ensuring that the bank's liabilities trade at par and that its monetary equilibrium is unique.

The classical discipline on over-issuance operates precisely through this quality channel.

If a bank expands note issuance relative to its reserve holdings, adverse clearings at the clearinghouse induce an outflow of gold. Because gold carries $\tilde{q} = 1$, its depletion mathematically reduces the average fundamental quality $\tilde{q}_t^b$ of the remaining asset portfolio. Absent balance-sheet contraction, continued reserve losses will drive $\tilde{q}_t^b$ below $\tilde{q}^*$. At that point, the liquidity premium on the bank's specific notes vanishes, the notes cease to circulate at par, and the individual producer exits. Monetary discipline thus arises endogenously from the interaction between reserve drains and the quality threshold structure of the model.

Remark 23 (Dollarized economies). In a fully dollarized economy, the domestic authority cannot act as a LOLR because it cannot create base money. To compensate, the private banking system must operate with substantially larger self-insurance buffers of highly liquid reserves (assets with $\tilde{q} = 1$). Thus, the banking system effectively maintains its aggregate portfolio quality $\tilde{q}_t$ securely above $\tilde{q}^{**}$, ensuring the fragile region is structurally avoided. This reliance on international reserves as quality enhancers explains why dollarized systems often exhibit significantly more conservative, liquidity-heavy balance sheets than comparable economies with domestic central banks.

6 Conclusion

This paper has proposed a theory of financial crises based on the concept of monetary production: the technology that transforms backing assets into liquid media accepted as money. The central result is a computable crisis threshold $\tilde{q}^* = (1 - \beta - \omega)/(\beta\kappa)$—derived by reinterpreting the Brunnermeier and Niepelt (2019) general liquidity constraint as a production technology—below which monetary production collapses through a saddle-node bifurcation.

The framework yields two main contributions.

1. A formal theory of monetary production with a computable threshold. I define monetary producers as agents whose instruments enter the liquidity constraints of households, and derive a production frontier that is increasing and linear in backing-asset quality. The threshold $\tilde{q}^*$ arises from the self-referential structure of the monetary equilibrium: liquid media are accepted because their issuer is solvent, and the issuer is solvent partly because its liquid media are accepted. The threshold is computable from four primitives (β, λ, υ, κ) and is invariant to institutional form. A quantitative calibration demonstrates the three-region equilibrium structure and shows that $\tilde{q}^*$ and the fragility index $\mathcal{G} = \tilde{q}^{**} - \tilde{q}^*$ respond to each primitive in the direction predicted by the theory.

2. Unification of crisis types. Banking crises, currency crises, sovereign debt crises, and debasement crises are all instances of the same mechanism operating on different sets of backing assets. The "this time is different" syndrome of Reinhart and Rogoff (2009)—the recurring belief that new financial innovations have eliminated crisis risk—is the belief that $\tilde{q}^*$ has shifted, when in fact $\tilde{q}^*$ depends on slow-moving primitives that are approximately constant across eras.

Limitations and future work. The closed-form results rely on a tractable CIA specification with log utility. While the qualitative results extend to the general BN framework (as argued in Section 2.8), a formal treatment under general preferences and liquidity constraints remains for future work. The empirical predictions—in particular, that monetary contraction precedes and causes output decline through the liquidity channel, and that crises can occur without coordination failures—are testable using the Jordà et al. (2017) macrofinancial history database, FRED monthly data, and the Nagel (2016) liquidity premium series; future empirical work will establish rigorous causal identification. The policy implications—particularly for LOLR design, macroprudential regulation targeting $\tilde{q}_t$, and the free-banking debate—deserve dedicated analysis that accounts for political-economy constraints. Finally, a formal model of commodity coinage (the debasement case) with endogenous minting technology and Gresham's Law dynamics is a natural companion extension.

The deeper lesson is methodological. Monetary economics has long treated money creation as either a mechanical process (money multiplier) or a policy instrument (central bank balance sheet). By treating it as a production technology—subject to input quality and critical thresholds—I connect the theory of money to the theory of production, bringing the analytical tools of the latter to bear on the central question of the former: when and why does the monetary system fail?

References

Acharya, V. V., Philippon, T., Richardson, M., and Roubini, N. (2009). The financial crisis of 2007-2009: Causes and remedies. In *Restoring Financial Stability: How to Repair a Failed System*, pages 1–56. John Wiley & Sons.

Anari, A., Kolari, J., and Mason, J. (2005). Bank asset liquidation and the propagation of the u.s. great depression. *Journal of Money, Credit and Banking*, 37(4):753–773.

Bagehot, W. (1873). *Lombard Street: A Description of the Money Market*. Henry S. King, London.

Baron, M., Verner, E., and Xiong, W. (2021). Banking crises without panics. *Quarterly Journal of Economics*, 136(1):51–113.

Bernanke, B. S., Gertler, M., and Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. *Handbook of Macroeconomics*, 1:1341–1393.

Brunnermeier, M. K. and Niepelt, D. (2019). On the equivalence of private and public money. *Journal of Monetary Economics*, 106:27–41.

Brunnermeier, M. K. and Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421.

Brunnermeier, M. K. and Sannikov, Y. (2016). *The I Theory of Money*. NBER Working Paper No. 22533.

Caballero, R. J., Farhi, E., and Gourinchas, P.-O. (2017). The safe assets shortage conundrum. *Journal of Economic Perspectives*, 31(3):29–46. See also: "On the Reach of Monetary Policy," *AER*, 2018.

Calomiris, C. W. and Kahn, C. M. (1991). The role of demandable debt in structuring optimal banking arrangements. *American Economic Review*, 81(3):497–513.

Calomiris, C. W. and Mason, J. R. (2003). Fundamentals, panics, and bank distress during the depression. *American Economic Review*, 93(5):1615–1647.

Cochrane, J. H. (2009). Lessons from the financial crisis. *Regulation*, 32:34.

Cochrane, J. H. (2023). *The Fiscal Theory of the Price Level*. Princeton University Press.

Correia, S., Luck, S., and Verner, E. (2026a). Bank failures: The roles of solvency and liquidity. Working Paper 34853, National Bureau of Economic Research.

Correia, S., Luck, S., and Verner, E. (2026b). Bank runs with and without bank failure. *arXiv preprint arXiv:2601.20285*.

Dang, T. V., Gorton, G., and Holmström, B. (2012). Ignorance, debt, and financial crises. Working paper.

Diamond, D. W. and Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3):401–419.

Diamond, D. W. and Rajan, R. G. (2001). Liquidity risk, liquidity creation, and financial fragility: A theory of banking. *Journal of Political Economy*, 109(2):287–327.

Eichengreen, B. (1992). *Golden Fetters: The Gold Standard and the Great Depression, 1919-1939*. Oxford University Press.

Farhi, E. and Tirole, J. (2018). Deadly embrace: Sovereign and financial balance sheets doom loops. *The Review of Economic Studies*, 85(3):1781–1823.

Friedman, M. and Schwartz, A. J. (1963). *A Monetary History of the United States, 1867-1960*. Princeton University Press.

Goldgar, A. (2008). *Tulipmania: Money, Honor, and Knowledge in the Dutch Golden Age*. University of Chicago Press.

Gorton, G. and Metrick, A. (2012). Securitized banking and the run on repo. *Journal of Financial Economics*, 104(3):425–451.

Gorton, G. and Ordoñez, G. (2014). Collateral crises. *American Economic Review*, 104(2):343–378.

Gorton, G. B. (2010). *Slapped by the Invisible Hand: The Panic of 2007*. Oxford University Press.

Harl, K. W. (1996). *Coinage in the Roman Economy, 300 B.C. to A.D. 700*. Johns Hopkins University Press.

He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *American Economic Review*, 103(2):732–770.

Holmström, B. and Tirole, J. (1998). Private and public supply of liquidity. *Journal of Political Economy*, 106(1):1–40.

Jiang, E. X., Matvos, G., Piskorski, T., and Seru, A. (2024). Monetary tightening and U.S. bank fragility in 2023: Mark-to-market losses and uninsured depositor runs? *Journal of Financial Economics*, 159:103899.

Jordà, Ò., Schularick, M., and Taylor, A. M. (2017). Macrofinancial history and the new business cycle facts. *NBER Macroeconomics Annual*, 31(1):213–263.

Kaminsky, G. L. and Reinhart, C. M. (1999). The twin crises: The causes of banking and balance-of-payments problems. *American Economic Review*, 89(3):473–500.

Kindleberger, C. P. and Aliber, R. Z. (2005). *Manias, Panics, and Crashes: A History of Financial Crises*. John Wiley & Sons, 5th edition.

Kiyotaki, N. and Moore, J. (2012). Liquidity, business cycles, and monetary policy. *Journal of Political Economy*. Working paper, Princeton University.

Krugman, P. (1979). A model of balance-of-payments crises. *Journal of Money, Credit and Banking*, 11(3):311–325.

Lagos, R. (2006). Inside and outside money. Staff Report 374, Federal Reserve Bank of Minneapolis.

Lagos, R. and Wright, R. (2005). A unified framework for monetary theory and policy analysis. *Journal of Political Economy*, 113(3):463–484.

Lucas, R. E. (1984). Money in a theory of finance. *Carnegie-Rochester Conference Series on Public Policy*, 21:9–45.

McCallum, B. T. (1983). The role of overlapping-generations models in monetary economics. *Carnegie-Rochester Conference Series on Public Policy*, 18:9–44.

McLeay, M., Radia, A., and Thomas, R. (2014). Money creation in the modern economy. *Bank of England Quarterly Bulletin*, 54(1):14–27.

Nagel, S. (2016). The liquidity premium of near-money assets. *Quarterly Journal of Economics*, 131(4):1927–1971.

Obstfeld, M. (1996). Models of currency crises with self-fulfilling features. *European Economic Review*, 40(3–5):1037–1047.

Reinhart, C. M. and Rogoff, K. S. (2009). *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press.

Sargent, T. J. and Velde, F. R. (2002). *The Big Problem of Small Change*. Princeton University Press.

Schnabel, I. (2004). The german twin crisis of 1931. *Journal of Economic History*, 64(3):822–871.

Selgin, G. A. (1988). *The Theory of Free Banking: Money Supply under Competitive Note Issue*. Rowman & Littlefield, Totowa, NJ.

Sidrauski, M. (1967). Rational choice and patterns of growth in a monetary economy. *American Economic Review*, 57(2):534–544.

Svensson, L. E. (1985). Money and asset prices in a cash-in-advance economy. *Journal of Political Economy*, 93(5):919–944.

Thompson, E. A. (2007). The tulipmania: Fact or artifact? *Public Choice*, 130(1-2):99–114.

White, L. H. (1984). *Free Banking in Britain: Theory, Experience, and Debate, 1800–1845*. Cambridge University Press, Cambridge.

Williamson, S. D. (2012). Liquidity, monetary policy, and the financial crisis: A new monetarist approach. *American Economic Review*, 102(6):2570–2605.

A Proofs and Derivations

A.1 Derivation of the Liquidity Kernel

I derive the liquidity kernel for the CIA constraint under the tractable specification (Theorem 2).

The household solves:

$$\max \mathbb{E}_0\left[\sum_{t=0}^{\infty} \beta^t \ln(c_t^h)\right] \quad \text{s.t.} \quad (3),\ (16)$$

The Lagrangian is:

$$\mathcal{L} = \sum_t \beta^t \Bigg\{ \ln(c_t^h) + \tilde{\mu}_t^h \left[\sum_j a_{t-1}^{j,h}(p_t^j + z_t^j) - \sum_j a_t^{j,h} p_t^j - c_t^h - \tau_t^h \right] + \tilde{\lambda}_t^h \left[a_t^{M,h} p_t^M + \sum_b a_t^{D^b,h} p_t^{D^b} v - c_t^h \right] \Bigg\}$$

Taking the first-order condition with respect to holdings of liquid media issued by producer i (e.g., deposits D^b or notes N^b) with velocity weight v. Under the CIA specification (16), the liquidity constraint at time t depends on current-period portfolio positions $a_t^{j,h} p_t^j$—not on inherited positions from the previous period. Consequently, the choice of $a_t^{j,h}$ affects L_t^h but not L_{t+1}^h:

$$-\tilde{\mu}_t^h p_t^j + \mathbb{E}_t\left[\tilde{\mu}_{t+1}^h (p_{t+1}^j + z_{t+1}^j)\right] + \tilde{\lambda}_t^h v\, p_t^j = 0$$

Dividing by $\tilde{\mu}_t^h p_t^j$ and defining $\lambda_t^h \equiv \tilde{\lambda}_t^h / \tilde{\mu}_t^h$, $\mu_{t,t+1}^h \equiv \tilde{\mu}_{t+1}^h / \tilde{\mu}_t^h$:

$$1 - \lambda_t^h v = \mathbb{E}_t\left[\mu_{t,t+1}^h \frac{p_{t+1}^j + z_{t+1}^j}{p_t^j}\right]$$

Rearranging:

$$p_t^j = \mathbb{E}_t\left[\mu_{t,t+1}^h \cdot \frac{1}{1 - \lambda_t^h v} \cdot (p_{t+1}^j + z_{t+1}^j)\right]$$

Defining the liquidity kernel:

$$\Lambda_{t,t+1}^{j,h} = \frac{1}{1 - \lambda_t^h v} \tag{47}$$

The absence of the forward-looking λ_{t+1}^h term—relative to the general-framework kernel of Brunnermeier and Niepelt (2019)—is a direct consequence of the CIA specification. The general liquidity constraint (4) allows inherited positions $a_{t-1}^{j,h}(p_t^j + z_t^j)$ to enter L_t^h, which would generate the BN kernel $(1 + \lambda_{t+1}^h v)/(1 - \lambda_t^h v)$ with both a forward and a contemporary shadow price. Under the CIA (16), only current positions $a_t^{j,h} p_t^j$ enter L_t^h, yielding the simpler form $1/(1 - \lambda_t^h v)$.

I obtain:

$$p_t^j = \mathbb{E}_t\Big[\mu_{t,t+1}^h \, \Lambda_{t,t+1}^{j,h} \, (p_{t+1}^j + z_{t+1}^j)\Big] \tag{48}$$

as required.

Remark 24 (Regularity Condition). The liquidity kernel $\Lambda_{t,t+1}^{j,h} = 1/(1 - \lambda_t^h v)$ is well-defined and strictly positive provided $\lambda_t^h v < 1$, i.e., the shadow value of the liquidity constraint times the velocity weight is strictly less than one. This condition holds in any interior equilibrium where the household's marginal utility of wealth exceeds the marginal value of relaxing the liquidity constraint. As $\lambda_t^h v \to 1$, the kernel diverges: $\Lambda \to \infty$. This divergence is the mathematical expression of the crisis boundary—the point at which the monetary equilibrium ceases to exist.

A.2 Derivation of the Steady-State Shadow Price

Under the CIA specification (16), the household chooses its portfolio $a_t^{j,h}$ at time t, and these holdings simultaneously enter the liquidity constraint at time t. The CIA binds as $\mathcal{M}_t/P_t \geq c_t$, where $\mathcal{M}_t$ represents the total liquidity value of the portfolio chosen at time t.

In steady state, the first-order condition for a riskless nominal bond with gross return $1 + i$ (which does not enter L_t^h, so $v = 0$) yields the Fisher equation:

$$1 = \beta(1 + i)\frac{P_t}{P_{t+1}}$$

For outside money M (with $v^M = 1$), the kernel pricing equation (48) gives:

$$p_t^M = \mathbb{E}_t\left[\frac{\beta}{1 - \lambda_t} \cdot p_{t+1}^M\right]$$

In steady state with $\lambda_t = \lambda$ constant and $P_{t+1}/P_t = 1 + \pi$, using $p_t^M = 1/P_t$:

$$1 - \lambda = \frac{\beta}{1 + \pi} = \frac{1}{1 + i}$$

where the last equality uses the Fisher relation. Solving:

$$\lambda = \frac{i}{1 + i} \tag{49}$$

This is the formula used throughout the main text. It satisfies $\lambda \in (0, 1)$ for $i > 0$, $\lambda \to 0$ as $i \to 0$ (the Friedman rule), and $\lambda \to 1$ as $i \to \infty$ (the regularity boundary).

A.3 Derivation of $\tilde{q}^*$ from Primitives

I derive the closed-form expression for $\tilde{q}^*$ from the three conditions: bank solvency, liquidity premium, and money market clearing.

Step 1: Solvency condition. At the solvency boundary $E = 0$:

$$A = \mathcal{M}, \quad \text{i.e.,} \quad F + LP = \mathcal{M}$$

Step 2: Exogenous quality. The quality parameter $\tilde{q} \in (0, 1]$ is exogenous: it determines the fraction of promised cash flows actually delivered by the backing assets. Fundamental value is:

$$F = \frac{\tilde{q} \cdot y^A}{r}$$

where $r = 1/\beta - 1$ is the steady-state real interest rate. When $\tilde{q} = 1$, the asset is fully performing; when $\tilde{q} < 1$, the asset delivers only a fraction $\tilde{q}$ of its promised cash flow each period.

Step 3: Liquidity premium. The backing assets are perpetuities with real cash flow y^A. In steady state with kernel Λ:

$$A = \frac{\beta \Lambda y^A}{1 - \beta\Lambda}, \qquad F = \frac{\beta y^A}{1 - \beta}$$

so the liquidity premium is:

$$LP = A - F = \beta y^A \left[\frac{\Lambda}{1 - \beta\Lambda} - \frac{1}{1 - \beta} \right] = \frac{\beta y^A (\Lambda - 1)}{(1 - \beta\Lambda)(1 - \beta)}$$

Step 4: Kernel expansion. With $\omega = \lambda v$ and $\Lambda = 1/(1 - \omega)$:

$$\Lambda - 1 = \frac{\omega}{1 - \omega}, \qquad 1 - \beta\Lambda = 1 - \frac{\beta}{1 - \omega} = \frac{1 - \omega - \beta}{1 - \omega}$$

Step 5: Substitution. The liquidity premium on a backing asset with cash flow z becomes:

$$LP = \frac{\beta z \cdot \frac{\omega}{1-\omega}}{\frac{1-\omega-\beta}{1-\omega} \cdot (1 - \beta)} = \frac{\beta \omega \, z}{(1 - \omega - \beta)(1 - \beta)}$$

Define the **liquidity production coefficient**:

$$\Phi(\omega, \beta) \equiv \frac{\beta\omega}{(1-\omega-\beta)(1-\beta)} \tag{50}$$

so that $LP = \Phi \cdot z$ for an asset with cash flow z. For a backing asset with quality $\tilde{q}$ delivering cash flow $z = \tilde{q} \cdot y^A$, the liquidity premium is $LP = \Phi \cdot \tilde{q} \cdot y^A$. The coefficient Φ is well-defined and positive under the parameter restriction $\omega < 1 - \beta$.

Step 6: Money market clearing. $\mathcal{M}/P = Y$, so $\mathcal{M} = PY$ in nominal terms, and in real terms $\mathcal{M}/P = Y$.

Step 7: Solvency at the boundary. From Steps 1–2 and 5, at $E = 0$:

$$F + LP = Y$$

Substituting $F = \tilde{q} \cdot y^A/r$ and $LP = \Phi \cdot \tilde{q} \cdot y^A$:

$$\tilde{q} \cdot y^A \left(\frac{1}{r} + \Phi \right) = Y$$

Note that $1/r + \Phi = \beta/(1-\beta) + \beta\omega/[(1-\omega-\beta)(1-\beta)] = \beta/(1-\omega-\beta)$.

Step 8: Solve for $\tilde{q}^*$. Defining $\kappa = y^A/Y$:

$$\tilde{q} \cdot \kappa \cdot \frac{\beta}{1-\omega-\beta} = 1$$

$$\boxed{\tilde{q}^* = \frac{1-\beta-\omega}{\beta\kappa}} \tag{51}$$

This can equivalently be written as $\tilde{q}^* = 1/[\kappa(1/r + \Phi)]$. For the existence of a monetary equilibrium ($\tilde{q}^* < 1$), one needs $\kappa > (1-\beta-\omega)/\beta$, i.e., the backing-asset cash-flow base must be large enough relative to the economy's liquidity needs. The upper threshold (where the LP is not required for solvency) is $\tilde{q}^{**} = r/\kappa = (1-\beta)/(\beta\kappa)$. The fragility index is:

$$\mathcal{G} = \tilde{q}^{**} - \tilde{q}^* = \frac{\omega}{\beta\kappa} \tag{52}$$

Note that $\mathcal{G}$ can also be written as $r\Phi$, since $r\Phi = [(1-\beta)/\beta] \cdot [\beta\omega/((1-\omega-\beta)(1-\beta))] = \omega/(1-\omega-\beta)$. However, $\omega/(\beta\kappa)$ is the more direct expression from the threshold formulas.

A.4 Comparative Statics of $\tilde{q}^*$

With respect to ω. From $\tilde{q}^* = (1-\beta-\omega)/(\beta\kappa)$:

$$\frac{\partial \tilde{q}^*}{\partial \omega} = -\frac{1}{\beta\kappa} < 0$$

Higher effective liquidity intensity mathematically *lowers* the deterministic crisis threshold $\tilde{q}^*$: a higher ω generates a larger liquidity premium (Φ), providing a market-value cushion that sustains solvency at lower fundamental quality. However, this solvency relies on endogenous confidence. The fragility index $\mathcal{G} = \tilde{q}^{**} - \tilde{q}^* = \omega/(\beta\kappa)$ strictly increases:

$$\frac{\partial \mathcal{G}}{\partial \omega} = \frac{1}{\beta\kappa} > 0$$

Thus, economies with higher ω are systematically more fragile: the fragile region is wider, amplification within it is stronger, and susceptibility to crossing $\tilde{q}^*$ via sunspot-driven multiple equilibria increases.

For reference, the derivative of Φ with respect to ω is:

$$\frac{\partial \Phi}{\partial \omega} = \frac{\beta}{(1-\omega-\beta)^2} > 0$$

This follows from the quotient rule applied to $\Phi = \beta\omega/[(1-\omega-\beta)(1-\beta)]$, noting that $\partial[(1-\omega-\beta)(1-\beta)]/\partial\omega = -(1-\beta)$ and $(1-\omega-\beta)+\omega = 1-\beta$.

With respect to κ. From $\tilde{q}^* = (1-\beta-\omega)/(\beta\kappa)$:

$$\frac{\partial \tilde{q}^*}{\partial \kappa} = -\frac{1-\beta-\omega}{\beta\kappa^2} = -\frac{\tilde{q}^*}{\kappa} < 0$$

Higher κ (backing assets generate more output) lowers $\tilde{q}^*$, making crisis less likely. Intuitively, more productive backing assets generate a larger fundamental value per unit, so the balance sheet can absorb larger quality shocks before breaching the solvency boundary. The fragility index also responds: $\partial\mathcal{G}/\partial\kappa = -\omega/(\beta\kappa^2) < 0$, so the fragile region *narrows* with more productive backing.

With respect to β. From $\tilde{q}^* = (1-\beta-\omega)/(\beta\kappa)$:

$$\frac{\partial \tilde{q}^*}{\partial \beta} = \frac{-\beta\kappa - (1-\beta-\omega)\kappa}{(\beta\kappa)^2} = \frac{-(1-\omega)}{\beta^2\kappa} < 0$$

since $\omega < 1$. More patient economies have *strictly lower* crisis thresholds—they are more resilient to backing-asset deterioration. The sign is unambiguous under the parameter restriction $\omega < 1 - \beta$.

For reference, the derivative of Φ with respect to β is:

$$\frac{\partial \Phi}{\partial \beta} = \frac{\omega(1 + \beta - \omega - 2\beta^2)}{(1 - \omega - \beta)^2(1 - \beta)^2}$$

which is positive for typical calibrations ($\beta \approx 0.96$, ω small) since $1 + \beta - 2\beta^2 > 0$ for $\beta \in (0, 1)$.

A.5 Robustness: Money-in-Utility Specification

Under a money-in-utility (MIU) specification, the household maximizes:

$$\mathbb{E}_0\left[\sum_{t=0}^{\infty} \beta^t \Big[u(c_t^h) + \xi(m_t^h)\Big]\right]$$

where $m_t^h = (a_t^{M,h} p_t^M + \sum_b a_t^{D^b,h} p_t^{D^b} v)/P_t$ is real money holdings and ξ is increasing and concave.

The key difference from CIA is that the liquidity motive operates through a first-order condition:

$$\xi'(m_t^h) = \lambda_t^h$$

rather than a binding constraint. Since ξ' is continuous and strictly decreasing, the shadow cost of liquidity λ_t^h varies continuously with real money holdings, rather than jumping discretely between binding and non-binding.

The liquidity kernel under MIU takes the form:

$$\Lambda^{\text{MIU}} = \frac{1}{1 - \lambda_t v}$$

where $\lambda_t = \xi'(m_t)/u'(c_t)$ is the ratio of the marginal utility of liquidity to the marginal utility of consumption. Because the CIA specification depends only on current-period positions, both the CIA and MIU kernels take the same reduced form $1/(1 - \lambda_t v)$. The MIU kernel lacks the forward-looking $(1 + \lambda_{t+1} v)$ numerator found in the general BN framework kernel—precisely because the MIU liquidity motive, like the CIA constraint, operates through contemporaneous holdings rather than through inherited positions that would generate a forward-looking shadow-price term. Qualitatively, however, $\Lambda^{\text{MIU}} \geq 1$ whenever $\lambda_t v > 0$, generating the same liquidity premium structure. The three-region structure of Theorem 7 is

preserved qualitatively:

- The monetary equilibrium exists for $\tilde{q} > \tilde{q}^*$ (analogous to $\tilde{q}^*$).
- Multiple equilibria exist for $\tilde{q}^* < \tilde{q} < \tilde{q}^{**}$.
- Only the autarkic equilibrium exists for $\tilde{q} < \tilde{q}^*$.

The thresholds $\tilde{q}^*$ and $\tilde{q}^{**}$ do not have closed-form expressions (they depend on the functional form of ξ), but the qualitative structure—and in particular the self-referential fixed-point mechanism that generates the crisis threshold—is identical. The CIA specification is therefore a tractable special case, not a knife-edge assumption.

A.6 Three-Component Price Decomposition

Iterating the Euler equation (48) forward T periods:

$$p_t^j = \mathbb{E}_t\left[\sum_{s=1}^{T} \mu_{t,t+s}^h \Lambda_{t,t+s}^{j,h} z_{t+s}^j\right] + \mathbb{E}_t\left[\mu_{t,t+T}^h \Lambda_{t,t+T}^{j,h} p_{t+T}^j\right]$$

where $\mu_{t,t+s}^h = \prod_{\tau=t}^{t+s-1} \mu_{\tau,\tau+1}^h$ and $\Lambda_{t,t+s}^{j,h} = \prod_{\tau=t}^{t+s-1} \Lambda_{\tau,\tau+1}^{j,h}$.

Adding and subtracting the fundamental value:

$$p_t^j = \underbrace{\mathbb{E}_t\left[\sum_{s=1}^{T} \mu_{t,t+s}^h z_{t+s}^j\right]}_{\text{Fundamental value (truncated)}} + \underbrace{\mathbb{E}_t\left[\sum_{s=1}^{T} \mu_{t,t+s}^h (\Lambda_{t,t+s}^{j,h} - 1) z_{t+s}^j\right]}_{\text{Liquidity premium (truncated)}} + \underbrace{\mathbb{E}_t\left[\mu_{t,t+T}^h \Lambda_{t,t+T}^{j,h} p_{t+T}^j\right]}_{\text{Residual}}$$

Taking $T \to \infty$ under appropriate transversality conditions yields equation (13). The bubble component $B_t^j = \lim_{T\to\infty} \mathbb{E}_t[\mu_{t,t+T}^h \Lambda_{t,t+T}^{j,h} p_{t+T}^j]$ is zero for assets with finite maturity and for infinitely-lived assets under the standard no-bubble condition augmented by the liquidity kernel.

For backing assets, I assume $B_t^j = 0$, so $A_t = F_t + LP_t$ as used throughout Section 3.