OESTERREICHISCHE NATIONALBANK
EUROSYSTEM

Working Paper 280

Stablecoins, Inflation and the Settlement Hierarchy

Martin Summer

ONB OESTERREICHISCHE NATIONALBANK EUROSYSTEM

OeNB Working Paper Series
84 pages

Stablecoins, Inflation and the Settlement Hierarchy

Would widely used stablecoins raise prices? This paper studies coins backed by government debt in a general equilibrium model of money and inflation. Such backing turns government bonds into spendable payment capacity, so the price level rises as adoption spreads. Monetary policy can offset the rise, but only at a real cost. And if private tokens were ever accepted in final settlement, the price level would lose its anchor. Regulation, not technology, decides these outcomes.

Author
Martin Summer
Oesterreichische Nationalbank,
Economic Research Division,
martin.summer@oenb.at

JEL classification
E31, E42, D52, E44, G21, G28

Keywords
Stablecoins, inflation, public and private money, monetary architecture

A channel of inflationary finance

Stablecoins backed by Treasury bills turn government debt into a payment instrument: the existing bond stock held as reserves becomes spendable, and newly issued debt arrives partly spendable. Both effects raise the price level, and they grow with the breadth of adoption.

Bounded, and a policy choice

Monetary policy can offset the price-level effect, but only by suppressing trade financed by bank credit. The channel's size is set by regulation: reserve rules steering issuers into government bills switch it on; backing in central bank money switches it off.

Final settlement must stay public

Current regulatory design leaves open an escalation path on which stablecoins become a competing money and nominal determinacy is at risk. The safeguard is constitutional, not prudential: taxes and bank obligations must remain settled in public money.

Stablecoins, Inflation and the Settlement Hierarchy

Martin Summer*

September 1, 2026

Abstract

Would widely adopted stablecoins affect the price level, and through which channels? We analyse this question in the monetary general-equilibrium framework of Dubey and Geanakoplos, in which payment capacity—money, bank credit and the spendable share of government bonds—determines nominal outcomes. Stablecoins backed by government debt open a channel of inflationary finance: the existing bond stock becomes spendable (a stock effect) and newly issued debt arrives spendable (a flow effect), with the price level rising in the payment capacity so created and in the breadth of adoption. Monetary policy can offset the rise, but only at the cost of suppressing the trade that bank credit finances. Beyond these bounded effects, current regulatory design does not close an escalation path on which stablecoins become a competing outside money, putting nominal determinacy at risk. Whether these effects materialise is largely a policy choice: they require adoption in ordinary payments, which regulation currently shapes.

Keywords: Stablecoins, inflation, public and private money, monetary architecture.

JEL: E31, E42, D52, E44, G21, G28

---

*Martin Summer, Economic Research Division, Oesterreichische Nationalbank, Otto Wagner Platz 3, POB-61, 1011 Vienna. Email: martin.summer@oenb.at. Acknowledgements: I thank Ulrich Bindseil, Johannes Binswanger, Rainer Böhme, Helmut Elsinger, Hans Gersbach, Michel Habib, Bernhard Haslhofer, Markus Knell, Loretta Mester, Dirk Niepelt, Paul Pichler, Nikola Tarashev, Branko Urošević and Beat Weber as well as an anonymous referee and seminar participants at OeNB and participants at the EEA meetings 2026 in Dublin for helpful comments. The views expressed in this paper do not necessarily reflect those of the Oesterreichische Nationalbank. *Statement on the use of AI tools:* This paper was prepared with the assistance of Claude (Anthropic). The text was written by the author; Claude was used under the author's direction for research and editing assistance—checking derivations and numerical results, verifying references and factual claims against primary sources, advising on structure, and drafting a limited number of passages, all of which the author reviewed. The R code of the companion package `dgme` was written by Claude under the author's direction and review, as documented in the package itself. The author takes full responsibility for the content of the paper.

Non-Technical Summary

Stablecoins are privately issued digital tokens designed to trade at par with an official currency—in practice almost always the US dollar—and backed largely by short-term government debt. They have grown into a market of several hundred billion dollars, and new legislation on both sides of the Atlantic—the US GENIUS Act and the EU's MiCA regulation—now decides how far they may move from crypto trading into ordinary payments. This paper asks what such a migration would mean for the price level and for the monetary order: the division of roles between public and private money.

The analysis is built on a general equilibrium model of money and inflation developed by Pradeep Dubey and John Geanakoplos, in which the price level is anchored by public money and moves with the economy's payment capacity: the money people hold, the credit banks extend, and—decisively for our question—the part of government debt that can effectively be spent. That spendable government debt raises the price level is a cornerstone of their analysis. The contribution of this paper is to trace stablecoins through their framework: coins backed by Treasury bills are exactly the financial engineering that turns government debt into a payment instrument.

The first finding is that widely adopted stablecoins open a channel of inflationary finance—the classical term for financing public expenditure with liabilities that come to function as means of payment; in the model the effect is a one-off rise in the price level, not a permanently higher inflation rate. The existing stock of bonds held as reserves becomes spendable, and newly issued debt arrives partly spendable, because stablecoin demand lowers the government's funding costs. Both effects raise the price level; they remain bounded as long as the monetary order stays intact, and monetary policy can offset them—at the cost of forgoing gains from trade. Their size is decided less by technology than by regulation: reserve rules steering issuers into government bills switch the channel on; backing in central bank money switches it off.

The second finding concerns a risk of a different category. The current regulatory design does not close an escalation path on which stablecoins stop being redeemable claims within the monetary order and begin to circulate as competing money. Were private tokens accepted in final settlement—for taxes and bank obligations—the economy would have two competing monies and the price level would lose its anchor. Guarding against this is not a matter of prudential fine-tuning but of the constitutional rules of the monetary order: if the price level is to keep its anchor, final settlement must remain in public money.

The policy conclusion is a division of labour. In the payment layer, private initiative is legitimate and conduct regulation—reserve composition, redemption terms, disclosure—is sufficient. The settlement layer is a natural public monopoly and must be protected. And the analysis provides little support for policies that would actively amplify demand for bond-backed coins: legislation of that kind makes the fiscal channel a feature rather than a side effect. For Europe, the framework speaks to the design of euro stablecoins, to the digital euro, and to the euro area's exposure to dollar coins.

Contents

1. Introduction

Digital payments saw a wave of innovations over the past decade. Among them stablecoins — crypto-tokens pegged to a sovereign currency and redeemable at par — have attracted much attention in policy circles and the payments industry. In the discussion stablecoins are often framed as the private-sector-led path to innovations in payments digitalisation and as an attractive alternative to public-sector-led initiatives to develop a central bank digital currency (CBDC).

This framing is recent. It is the latest stage of a development that began in 2019, when the social media company Facebook announced Libra (later rebranded as Diem), a permissioned, blockchain-based stablecoin payment system (Libra Association, 2019). The announcement spurred a wave of public-sector CBDC projects (Boar and Wehrli, 2021). Libra/Diem was finally abandoned due to regulatory opposition, and with its demise many retail CBDC projects lost momentum and were shelved, the ECB's digital euro project being the most prominent exception. Stablecoins, however, kept thriving: issuance grew from niche scale to a market within the crypto sector in the hundreds of billions of US dollars, concentrated in US-dollar-pegged coins from two issuers (BIS, 2025). The regulatory frameworks of MiCA in the EU (European Parliament and Council of the European Union, 2023) and the GENIUS Act in the US (U.S. Congress, 2025) legislated new rules potentially allowing a broader integration of stablecoins into the financial system. This created a public climate in which building on stablecoins for the future of payment innovation increasingly appeared as an alternative to the CBDC projects led by central banks.

In this paper we confront this framing with a systematic analysis based on economic theory. We ask what widespread adoption of stablecoins as a general payment instrument would mean macroeconomically and for the current monetary order more broadly. In its precise form, the question is: do stablecoins affect the price level — and if so, through which channels, and what determines the size of each channel?

The answer we arrive at can be summarised in two findings. The first is that widely adopted stablecoins open a channel of inflationary finance:[1] their financial engineering turns government debt into payment capacity — the existing stock of bonds becomes spendable now, and newly issued debt arrives spendable — and the price level rises with the capacity so created. The second is a risk of nominal de-anchoring: the current

[1]"Inflationary finance" is used throughout in its standard public-finance sense: the financing of public expenditure through liabilities that function as payment capacity—here, existing debt that becomes spendable and new debt that arrives spendable. In the one-period model the effect is a higher equilibrium price level—a statement about the level of prices, not their rate of change. The footnote at the opening of Section 3 states the exact sense, which follows Dubey and Geanakoplos (2026).

regulatory design does not close an escalation path on which stablecoins cease to be payment instruments within the existing monetary order and begin to circulate as a competing outside money — putting at risk the nominal anchor of the very currency they were designed to track. The two findings are not of the same kind and should not be weighed on the same scale: the first is a bounded effect within a monetary order that remains intact; the second concerns the survival of that order.

How these findings should be read depends on a fact worth stating at the outset: the price-level effects we identify do not follow from the existence of stablecoins alone. A stablecoin moves the price level only to the extent that it is actually used to pay for ordinary goods and services — and today stablecoins circulate almost entirely within crypto-asset markets. Payments are, moreover, a two-sided market: financial engineering can supply an instrument fit to pay, but paying for goods requires sellers equipped and willing to accept it. Whether this acceptance infrastructure gets built is not a natural constant. It is shaped by legislation, regulation and official endorsement — by precisely the policy climate just described.

We therefore do not read the inflationary push analysed in this paper as a forecast. It is in fact a policy choice. If policy leads stablecoins out of their crypto niche and into everyday payments, these price-level effects are part of what that decision buys; if it does not, they remain largely dormant. The paper is accordingly best read as a decision aid for a choice that is currently being made, in the legislative frameworks just described, and whose monetary consequences have so far received little systematic analysis.

In our analysis we use a model due to Dubey and Geanakoplos, recently published as the 2024 Klein Lecture (Dubey and Geanakoplos, 2026).[2] The choice of this particular model is dictated by our question: a stablecoin is, whatever else it may be, a new payment instrument, so any macroeconomic effect it has must operate through its impact on changes in what can be used to pay — a margin from which more standard monetary macromodels deliberately abstract, because they are designed to answer different questions. The Dubey–Geanakoplos model provides the language through which stablecoins can be understood from a monetary perspective. We use their model as our workhorse. All we add with respect to the model is some simplifications which help to focus on the specific

[2]The lecture consolidates a research programme spanning three decades, which most of the modern macroeconomic literature has taken little notice of. The finite-horizon general equilibrium model of money with banks first appears in Dubey and Geanakoplos (1992); Dubey and Geanakoplos (2003) develop the gains-to-trade analysis and the relation to the classical IS-LM model; Dubey and Geanakoplos (2006) treat determinacy with nominal assets and outside money; Dubey and Geanakoplos (2010) add credit cards as an example how changes in payment technology can affect the price level. The authors survey this lineage themselves in §1.11 of the published lecture. An earlier version of this paper built on the working-paper version (Dubey and Geanakoplos, 2024).

challenges of stablecoins.[3] Since the framework will be unfamiliar to many readers, we provide a detailed exposition in Section 2. Readers familiar with it can skim that section for the notation used here and jump right to Section 3, where the contributions of this paper begin.

An overview of the main results All results of this paper can be read from a single object: the economy's payment capacity and the settlement obligations behind it, developed in Section 3. Having one such object makes the competing claims of the stablecoin debate commensurable. Whether a fully backed coin is monetarily harmless, for instance, stops being a matter of framing: it depends on what the backing does to payment capacity, and the answer differs by reserve asset. The channels that follow form a taxonomy within this one framework, and the policy discussion of Sections 5 and 6 is organised by it.

Imagine a stablecoin issuer that takes in dollars and buys T-bills with the proceeds. It hands out coins that pay at the checkout. In this circuit the T-bills—promises of public money in the future—have become spendable today. And the dollars the issuer took in do not leave circulation: they move on to whoever sold the bills, while the coins spend on top of them. No party anywhere in the chain gave up more than it received; yet the economy as a whole can now pay for more, and the price level rises with it.

The paper calls this price-level effect the stock channel; Section 3.1 is its home. The result itself is Dubey and Geanakoplos' own: Theorem 4 of the Klein Lecture, under the title "When Bonds Become Money, Prices Rise" (Dubey and Geanakoplos, 2026). Proposition 1 restates it in our simplified setting and gives it institutional content; Corollary 1 adds what does not move: settlement absorption. The proposition shows that the price level responds to the moneyness of the bonds—the share of the bond stock that can actually be spent at par—and to nothing else about who holds them.

Why is this different from a money market fund? The large dollar stablecoins of today hold their reserves mainly in Treasury bills, and government money market funds have held much the same portfolio for decades and at far larger scale. The difference lies in the claim, not in the portfolio: the fund issues shares, and a share cannot pay at the checkout—it must first be turned back into bank money. The coin is built to pay. For a given reserve asset, then, what matters for the price level is what the claim written on it can do in payment, and Section 3.1 makes the fund sector the control case that identifies the channel.

Moneyness is therefore not a free parameter, and we do not treat it as one: it is produced by an institution, the stablecoin issuer, and we model that institution directly. Its balance sheet delivers moneyness as a function of two quantities a supervisor can read off published

[3] The Klein Lecture paper by Dubey and Geanakoplos (2026) as well as their previous monetary work is more general and geared towards the analysis of money and new monies more generally.

accounts—the scale of issuance and the composition of the reserve assets. Proposition 2 pins this function down. Making moneyness endogenous is the analytical core of the paper.

A second channel opens when the bond stock itself grows. This is not a stablecoin-specific mechanism: the comparative static exists in the economy without coins, and the coin enters only through its moneyness and its demand for bonds. What the stablecoin adds is amplification—newly issued debt arrives already spendable. Corollary 2 records this flow channel of fiscal amplification.

A sceptical reader may accept both channels and object that nobody buys groceries with stablecoins. The model agrees. Payment capacity rises only to the extent that coins are actually accepted in payment for goods, and at today's adoption—confined almost entirely to crypto-asset markets—the effective moneyness of the outstanding coins stays near zero. Section 3.4 makes this bound explicit. It is also why we read the channels as a lever rather than a forecast: whether the acceptance infrastructure gets built is the policy decision currently on the table, and that decision can unleash or keep dormant the inflationary pressure engineered into the existing coins.

This discount does not, however, dispose of our second finding, because the escalation it concerns does not need more adoption. It needs different things: holders who keep the coin instead of redeeming it, and debts written in the token itself. Neither requires the checkout, and crypto-asset markets show both today.

Section 4 traces where these requirements can lead, in three stages. Today the issuer collects dollars, invests them in T-bills and stands ready to redeem at par: the coin is a wrapped claim on public money. As holding replaces redemption and credit begins to be denominated in the token, the circuit becomes self-sustaining: settlement inside it no longer touches the sovereign unit. In the end state the token circulates as a second outside money beside the dollar. The current regulatory design does not close this path—the GENIUS Act's designation of T-bills as the effective canonical reserve asset aligns the Treasury's funding interest with the issuers'—and no stage requires anyone to intend de-anchoring: each step can look individually benign. What the end state means is precise: with two circulating outside monies the price level loses its anchor—it is indeterminate (Proposition 3). The model turns "losing the nominal anchor" from a metaphor into a theorem.

Sections 5 and 6 draw the policy consequences. As long as the monetary order remains anchored, the central bank can offset the price-level push, and the accounting says how much tightening is required; but the offset is costly in a precise sense: the bank tightens against payment capacity it did not create, and the burden falls on bank-financed trade (Section 5). Section 6 then organises the policy debate around one question: which monetary functions must be public, and which can be private? The frontier runs between payments and settlement: private initiative in the payment layer is bounded and can be contained; the settlement layer is a natural monopoly and must remain public. Three applications follow: a reading of MiCA and the GENIUS Act through this frontier (Section 6.1), a diagnosis of the European debate on euro stablecoins (Section 6.2), and a reframing of the CBDC question (Section 6.3).

We can now return to the frame this paper opened with: the popular comparison of a nimble private solution against a sluggish public project. Governor Waller (2021) gave it its canonical form in a speech that has since become famous: "After exploring many possible problems that a CBDC could solve, I am left with the conclusion that a CBDC remains a solution in search of a problem." The private sector, he argued, was already developing the payment alternatives a CBDC would merely duplicate. Our analysis shows this comparison to be a category error: what matters for monetary consequences is not the technology stack but the position in the settlement hierarchy—who is the carrier, what is the backing, who drives adoption. In that frame the comparison changes: the question is not private rails against a public project, but who carries the claim that travels on the rails. The price-level effects analysed in this paper attach to the carrier and the backing, not to the engineering; Section 6.3 spells out what this does, and does not, imply for the CBDC question.

Related research The rapidly growing literature on stablecoins is overwhelmingly concerned with financial stability: whether issuers can honour the promise of par-value redemption under stress. Anadu et al. (2024) establish a close functional parallel between stablecoins and money market funds, documenting that both are susceptible to runs triggered by small deviations from par. Ma et al. (2025) analyse how the restriction of direct redemption to a small set of intermediaries creates a two-tiered system in which retail holders bear secondary-market risk. Gorton and Zhang (2023) draw a historical parallel to the pre-Civil-War "wildcat banking" era. Aronoff et al. (2026) provide a detailed analysis of the operational plumbing on which par-value depends, showing that redemption requires Treasury liquidation through broker-dealers whose balance-sheet capacity is tightly constrained. On the theory side, Van Buggenum et al. (2024) study stablecoins traded in a secondary market in an overlapping-generations model, showing that redemption limits in the spirit of Diamond–Dybvig suspension of convertibility can stabilise rather than destabilise and providing a rationale for the interest prohibitions of MiCA and the GENIUS Act; in the same run-risk tradition, Li et al. (2026) cast the policy problem as a trade-off between stability and the incentive to issue, showing that unregulated issuers hold risky reserves and that safe-asset backing restores run-proofness when paired with other sources of issuer revenue; Gersbach and Zelzner (2025) study competition between issuers in rollover markets, decomposing the competitive margin into yields and issuance discounts and identifying a funds-concentration effect by which issuer exit reshapes the funding conditions of the surviving issuers. The regulatory literature reflects the same framing: Odinet and Tosato (2025) analyse MiCA and the GENIUS Act, and Gersbach et al. (2026) distil the stability perspective into a comprehensive policy agenda for stablecoins as private money.

Our paper asks a different and, we argue, prior question: what happens to the price level when the promise of par *is* kept? The stability literature takes the macroeconomic environment as given and asks whether issuers can meet redemptions. We take issuer solvency as given and ask what the reserve-management practices that sustain solvency imply for aggregate payment capacity and prices. The two perspectives are complementary: Aronoff et al. (2026) show what can go wrong with the promise of par; we show what happens to the economy when it holds. Asking the question this way also takes the industry at its word: the declared ambition of stablecoin issuers is a modern digital payment instrument, not a substitute banking system — and it is as payment instruments that the coins reach the price level.

The macroeconomic question has a long lineage but has not been applied to stablecoins. Gurley and Shaw (1960) argued that financial intermediaries create "near-monies" whose expansion loosens the grip of the money stock on nominal spending; Tobin (1963) pushed back on the grounds that portfolio substitution limits net expansion; Fama (1980) showed that price-level determinacy requires an outside-money anchor. The Dubey–Geanakoplos framework we employ, rooted in the strategic market games approach to money developed by Shubik (1999), formalises these insights in a general equilibrium setting (Dubey and Geanakoplos, 2026). In the monetary-hierarchy tradition of Mehrling (2020), our indeterminacy result (Section 4) makes precise what the settlement hierarchy accomplishes when stablecoins are redeemable claims on public money, and Hellwig (1985) provides the analysis of competition between outside monies that our three-stage escalation makes directly applicable to the stablecoin context.

Among the few genuinely monetary treatments of stablecoins, the closest to our concerns is Benigno (2026). He asks whether the institution that manages the unit of account must also manage the means of payment — a near relative of the frontier question that organises our Section 6 — and studies how administered rates and central-bank balance sheets govern liquidity premia and the price level when stablecoins compete with public liabilities in the provision of liquidity services. The difference between his analysis and ours lies in the margin that is modelled. In his framework the set of claims that provide liquidity services is a primitive of preferences, in the convenience-yield tradition of Krishnamurthy and Vissing-Jorgensen (2012), and policy varies their remuneration and backing; in ours the moneyness of an instrument is itself produced — by financial engineering and by adoption — and the object of analysis is what this production does to the price level and the nominal anchor. The two analyses thus divide the labour: remuneration design and backstop architecture on his side, the creation of payment capacity and its fiscal consequences on ours. The quantitative macroeconomic counterpart is Hofmann et al. (2026a), who embed stablecoins in a calibrated banking model and identify two opposing channels—a bank lending channel and a fiscal space channel, the latter the model-based analogue of our flow channel—with the balance between them decided by reserve regulation, the level of public debt and foreign demand.

A second monetary strand studies the international dimension. Azzimonti and Quadrini (2025) analyse how stablecoin reserve demand feeds the global appetite for US safe assets and the dollar's privilege, Ferrari Minesso and Siena (2026) the safe-asset channel through which dollar stablecoins tie foreign payment demand to US public debt, and Barthelemy et al. (2025) the consequences for the international monetary system; empirically, Ahmed and Aldasoro (2025) and Barthelemy et al. (2024) document the footprint of stablecoin reserve demand in T-bill yields and short-term funding markets, and Hofmann et al. (2026b) the dollarisation footprint of stablecoin inflows in emerging markets. These mechanisms operate through cross-border flows and asset prices; our flow channel operates through the domestic price level.

A third body of work does not address stablecoins directly but frames the questions they raise. Brunnermeier and Niepelt (2019) establish conditions under which private and public money are equivalent, and Niepelt (2026) organises the literature on central bank digital currency around the resulting irrelevance result; the statement in Section 6.3 that a CBDC moves no term of the payment-capacity accounting is the counterpart of that result in our framework. The design literature on central bank digital currency — Andolfatto (2021), Fernández-Villaverde et al. (2021), Keister and Sanches (2023) — studies what happens when the equivalence fails: deposit migration, bank funding and the intermediation consequences of a public option, which are precisely the two-tier effects our compression sets aside. Further back, Gorton and Pennacchi (1990) explain why privately produced liquidity takes the form of information-insensitive debt — the rationale behind the reserve-backed design of the coins we study — and Krishnamurthy and Vissing-Jorgensen (2012) measure the convenience yield on Treasury debt, establishing the tradition in which liquidity services are a primitive of preferences. What separates our paper from all of these strands is the margin none of them models: the literature either asks whether the promise of par can be kept, or takes the moneyness of the token as given. We take the promise as kept — and model the moneyness as produced.

2. The Dubey–Geanakoplos Model

2.1. The real side: An exchange economy

The real side of the economy is a pure exchange economy. There is a finite set $\mathcal{L} = \{1, \ldots, L\}$ of physical goods. A finite set $\mathcal{H}$ of households is characterised by endowments $e^h \in \mathbb{R}^L_+$ of the L goods and by preferences $u^h : \mathbb{R}^L_+ \to \mathbb{R}$ over consumption bundles, assumed to be continuous, strictly increasing, and concave, for every household $h \in \mathcal{H}$. We assume $e^h \neq 0$ for all $h \in \mathcal{H}$ and $\sum_{h \in \mathcal{H}} e^h \gg 0$. This description follows the usual conventions of general equilibrium theory and is deliberately sparse: the mechanisms studied in this paper operate through how transactions are settled, not through what is produced, and a pure exchange economy—as in Dubey–Geanakoplos—provides the minimal real structure on which the monetary architecture can be built. The real primitives contain no time: the endowment is a single bundle rather than a dated stream, and preferences rank consumption bundles, not consumption paths. Time enters the model only through the nominal or financial side, to which we now turn.

2.2. The nominal or financial side: Money and government bonds

The real side of the economy is static, but exchange is not synchronous. In any real exchange process, agents must pay for purchases before they receive the proceeds from their sales, so the Walrasian fiction of simultaneous net trade—in which expenditure is financed instantaneously by revenue—cannot hold. The financial side of the model exists to resolve this timing problem, and its structure is organised around it: there are two dates, $t = 0$ and $t = 1$. All goods trade takes place at $t = 0$; all financial obligations are settled at $t = 1$; consumption follows settlement. In the baseline model, all goods trade is settled in money (we extend the set of admissible payment instruments in Section 3). Money can take—as in our current monetary system—the form of cash or bank deposits. These are the two forms of money we distinguish in the baseline model. Purchasing power reaches households in three ways: as endowed money at $t = 0$, as delivery on government bonds at $t = 1$, and as bank credit at $t = 0$. The following paragraphs introduce each of these modes of money in turn.

Public or outside money. Household h starts at $t = 0$ with an endowment of $m^h \geq 0$ units of public or outside money, owned free and clear of debt. Both names record the same defining property: this money originates *outside* the private sector—issued by the *public* sector, the only agent that can create at will the very units in which its own liabilities are discharged—and no private agent owes an offsetting obligation against it, so that it constitutes net wealth to the household sector. Let $M \equiv \sum_{h \in \mathcal{H}} m^h$ denote the aggregate stock; we assume $M > 0$, so that the settlement medium exists in positive supply. Outside money is the unit of account: goods prices are quoted in it, and we write $p \in \mathbb{R}^L_{++}$ for the vector of money prices of goods at $t = 0$.[4] It is also the ultimate settlement asset: all payment obligations must eventually be discharged in outside money. Outside money is intrinsically valueless—it cannot be consumed and yields no direct utility. That it can nevertheless carry value in equilibrium rests on its institutional role: the rules of the economy *require* it as the medium in which all financial promises are denominated and all trades are settled. But the institutional role only creates the demand for the settlement medium; whether that demand supports a positive value is a further, economic question. In the one-period economy studied here, it does so precisely when the economy's gains from trade are large enough to carry the cost of transacting through money—a condition made precise in the existence discussion at the end of this section. How this argument fares beyond the one-period compression is taken up in the paragraph on the model's imagined dynamics below.

Government bonds. Household h also starts at $t = 0$ with an endowment $b^h \geq 0$ of government bonds, held free and clear. A government bond is a nominal promise: one bond delivers one unit of outside money at $t = 1$. Delivery occurs at settlement, where bond proceeds are netted against the household's other financial obligations. A bond is thus deferred outside money—a claim of the same kind as outside money, differing in the date at which it can pay, not in the nature of the promise. In the baseline model bonds cannot be used in goods trade at $t = 0$; whether and how they acquire that use is precisely the question of Section 3. Let $B \equiv \sum_{h \in \mathcal{H}} b^h \geq 0$ denote the aggregate stock; the special case $B = 0$ recovers the bond-free monetary economy with outside and inside money only, as studied in the authors' earlier work (Dubey and Geanakoplos, 1992; Dubey and Geanakoplos, 2003).

The endowment abstracts from a layer of institutional reality that deserves comment at first contact, because the paper leans on the endowment throughout. Few households hold government bonds outright. The household sector's ownership of the public debt runs almost entirely through intermediaries—money market and bond mutual funds,

[4]Dubey and Geanakoplos work with a pairwise price system—the price of object α in units of instrument β—because their market structure lets several instruments trade directly against goods at potentially unlinked prices. Our market structure never needs this generality: whenever a second instrument acquires a goods-market role in this paper, its price is tied to p. In Dubey and Geanakoplos' theorem on spendable bonds, goods bought against bonds trade at the deferred-settlement price—the money price marked up by the interest wedge. The par-redeemable stablecoin of this paper is different: it enters the payment constraint at face value, so coin and money buy goods at the same price p—a difference of substance, not of notation, taken up after Proposition 1 in Section 3.1. Coin prices in the two-money economy satisfy $p^S = \varepsilon p$ (Section 4.2). A single money price system therefore carries the entire analysis.

exchange-traded funds, pension claims—and what sits in a household's own account is a share in a vehicle, not the security. The endowment b^h is to be read accordingly, as *beneficial ownership*: the public debt ultimately owned by household h, through however many layers of intermediation. The reading involves the same kind of compression the model applies below to the monetary architecture, and it is licensed the same way: a fund share is a pass-through claim on the bonds behind it, and at settlement the chain nets out—the bond's delivery reaches its ultimate owner whether one intermediary stands in between or five. What intermediation does change is not who owns the debt but what the claim can do *before* settlement: a fund share can be redeemed into money, but it cannot itself pay for goods. In the baseline model this distinction is invisible, because there bonds cannot pay for goods in any form. It becomes the decisive distinction of Section 3.1, where the question is precisely which claims on the public debt the payment constraint accepts. There the money market fund—the largest of the intermediaries hidden behind the endowment b^h—reappears as the comparison case: it holds the same bonds, but its shares cannot pay for goods, and this is what makes the stablecoin's effect visible.

Private or inside money. Households can obtain additional purchasing power at $t = 0$ by borrowing from a banking sector. Banking is modelled as a linear technology $(-q, 1)$: it transforms one unit of future outside money into q units of current liquidity, where $q \in (0, 1)$ is the price of future outside money.[5] The household's position in this technology is described by a scalar $\ell^h \leq 0$: going short ($\ell^h < 0$) means taking a loan—the household receives $-q\ell^h > 0$ units of liquidity at $t = 0$ and must repay $-\ell^h > 0$ units of outside money at $t = 1$. Inside money is thus a claim on future outside money: bank credit expands the means of payment today but must be settled in outside money tomorrow.

The banking sector operating this technology is competitive. We adopt an interpretation that Dubey and Geanakoplos themselves offer for the model's credit arrangement: banks are licensed by the public sector to create money by lending, and pay a levy on each dollar so created to the public sector (Dubey and Geanakoplos, 2026). Free entry drives profit net of the levy to zero, so the terms set by the monetary authority are passed through to households one for one; banks accommodate whatever borrowing households choose at these terms. Their conduct in equilibrium is best understood as the zero-profit outcome that is forced by competition with a costless linear technology. Let $D \equiv -\sum_{h\in\mathcal{H}} \ell^h \geq 0$ denote aggregate credit—the total face value of inside money created. Since banks earn no margin, the interest wedge $(1 - q)D$ accrues in full to the public sector. This flow is not a technicality: it is the levy on money creation just described—Dubey and Geanakoplos speak of "the tax rate on the creation of money by government-licensed commercial

[5]Both bounds are economically meaningful: the Walrasian limit $q \to 1$ is discussed with the collapsed representation below, and the lower bound below which monetary equilibrium ceases to exist appears in the existence discussion at the end of this section.

banks”—and it will determine the accounting identity that closes the model at settlement. No household may default on its financial obligations.

A remark on instruments. Dubey and Geanakoplos take the quantity of new loans as the primary policy parameter, with the interest rate emerging in equilibrium, and note explicitly that the bank could equivalently set the rate and let the quantity emerge. We adopt this second, equivalent formulation throughout: q is the policy instrument, and aggregate credit D is endogenous.

Outside wealth. Define $W \equiv M + B$ and call it *outside wealth*: the total of claims on the public sector that households hold free and clear of any offsetting obligation. Dubey and Geanakoplos call this variable fiat wealth.[6] These claims arrived in households’ hands with no strings attached; an injection of endowment money is, as Dubey and Geanakoplos note, essentially what Milton Friedman called helicopter money.[7] Why an injection of this kind—rather than an ordinary central-bank operation—must stand at the beginning of the arrangement is taken up below, in the paragraph on the model’s imagined dynamics. The instruments of the model are classified by two distinctions: whether a claim is held free and clear or is offset by the holder’s own obligation, and whether it delivers purchasing power at $t = 0$ or at $t = 1$.

	delivers at $t = 0$	delivers at $t = 1$
held free and clear	outside money m^h	government bond b^h
offset by own obligation	inside money (bank credit)	loan repayment $-\ell^h$

The economically meaningful boundary runs between the rows, not the columns: the first row constitutes net wealth to the household sector and sums to W, while the entries of the second row are the two legs of a single contract and net to zero on the household’s own balance sheet. The columns differ only in timing.

Institutional interpretation. In modern monetary systems, outside money corresponds to liabilities of the central bank—both currency in circulation and reserve balances held by commercial banks. These liabilities enter circulation through the central bank’s asset purchases, typically of high-quality debt instruments of the public or private sector. The terms at which the central bank acquires these claims—the interest rate it charges, the collateral it demands—are the operational content of monetary policy. Government bonds

[6]We rename the variable because the everyday connotations of “fiat” invite readings the analysis neither needs nor makes; nothing of substance depends on the label.

[7]The baseline model has no taxes. Where the government levies lump-sum taxes, W is to be understood net of them: outside wealth is the unbacked residual of public liabilities over the private sector’s obligations to the public sector, not the gross stock of those liabilities.

correspond to the debt of the treasury. In the model the treasury and the central bank are consolidated into a single public sector: the household sector then faces two kinds of public liabilities, one that settles obligations now and one that becomes settlement money later.[8] The resulting architecture is *two-tiered*: the central bank provides liquidity to commercial banks, which in turn provide liquidity to households and firms.

The model deliberately compresses this architecture. Our abstraction consolidates currency and reserves into a single stock of outside money $(m^h)_{h \in \mathcal{H}}$, held directly by households, and replaces the two-tiered intermediation chain with a competitive banking sector that stands between households and the settlement mechanism. The price q reflects the terms on which the monetary authority licenses the creation of liquidity by the banking sector, which competition passes through to households without margin; the implied interest rate is $r = 1/q - 1$. This compression is justified because our analysis concerns the *aggregate payment capacity* of the economy—the total liquidity available to finance goods transactions—rather than the internal plumbing through which central bank liquidity reaches end users. For this question, what matters is not whether a unit of outside money sits in a household's wallet or in a bank's reserve account, but the total quantity in circulation and the credit terms at which it can be leveraged. Both are preserved in the reduced model: the exogenous stock M captures the aggregate supply of outside money, and q captures the terms on which the monetary authority allows the private sector to expand purchasing power beyond that stock. The key economic feature is captured by the abstraction: payment at the time of trade requires access to a scarce settlement medium that is not created inside the private sector—and the economy's endowed claims on that sector's counterpart, the public sector, come in two maturities, only one of which can pay at the time of trade.

The one-period model and its imagined dynamics. Readers accustomed to dynamic models will want to know where the one-period compression cuts, and the answer can be given once and for all: a finite monetary economy has a special beginning, a special end, and an interim that works differently from either. At the beginning, claims on the public sector must be placed into private hands free and clear; at the end, the entire outside wealth returns to the public sector through the settlement absorption condition derived below; in between—in any imagined dynamic extension—money is created and extinguished only through balance-sheet operations that swap equal values. The endowed stocks M and B are the point at which this structure is most often felt to be artificial, and the right place to

[8] This consolidation is a good approximation for the United States, on which our institutional discussion focuses. The focus is not arbitrary: the economically significant stablecoin markets are dollar-denominated, their issuers hold dollar reserve assets, and their regulatory treatment under the GENIUS Act – while designed for the US jurisdiction– has consequences and reach of global concern, not least for Europe. The euro area, where the consolidation is neither legally nor structurally available, is taken up in Section 6.2.

explain it. In reality nobody is endowed with money: outside money enters circulation through the balance-sheet operations described above, in which the central bank issues its liabilities only against assets of equal value. The model's accounting shows exactly what this observation gets right and what it misses. Every balance-sheet operation is a swap—money against bonds, or money against a repayment promise. A swap exchanges equal values: it can change the composition of the private sector's claims and, where a loan is involved, create matching entries in the second row of the classification table, but no sequence of swaps can leave the private sector holding positive net claims on the public sector. Monetary equilibrium, on the other hand, requires precisely such claims: at settlement the private sector surrenders its entire outside wealth as interest on inside money—the accounting identity anticipated in the discussion of the banking sector and derived below. With $W = 0$ there is nothing to surrender, no reason to lend or borrow, and no monetary trade. Once the system is running, offsetting claims are all that is needed: operations of the ordinary kind create and extinguish money against debt, and every unit so created foreshadows its own exit. But that the system be *started* requires a non-swap event—an initial transfer of claims to households, held free and clear. Nothing confines such events to the start: in a dynamic extension the public sector can add to outside wealth at any date by further expenditure not covered by taxes; what is structurally necessary is only that the *first* injection be of this kind. The helicopter drop is thus not an expository shortcut but a structural necessity, and the endowments (m^h, b^h) are best read as compressed fiscal history: the sediment of all past non-swap events, whenever they occurred—public expenditure paid for with liabilities that were never taxed back.[9]

2.3. Net trades and the settlement constraints

Given the monetary and financial instruments, we now characterise the constraints governing their use in trade. Let $\tau^h \equiv x^h - e^h$ denote household h's net trade in goods, where $x^h \in \mathbb{R}^L_+$ is the post-trade consumption bundle, and decompose it into net purchases $\tau^{h,+} \equiv \max(\tau^h, 0)$ and net sales $\tau^{h,-} \equiv \max(-\tau^h, 0)$, so that $\tau^h = \tau^{h,+} - \tau^{h,-}$.

Households can also trade bonds against money at $t = 0$. Let β^h denote household h's net bond purchase, subject to $\beta^h \geq -b^h$: no household can sell more bonds than it owns. Note that the bond is a claim to one unit of outside money at $t = 1$, and the loan market already prices such a claim exactly at q. Both must therefore trade at the same price; otherwise there would be an arbitrage opportunity.

[9] This reading takes a deliberate stand on the classic question of Barro (1974), whether government bonds are net wealth: by construction, the endowed bonds comprise only the part of the public debt that was never taxed back, not the outstanding stock in general. I have to thank Markus Knell as well as Beat Weber for directly and indirectly forcing me to better explain the nature of the one period model in relation to the richer (imagined) dynamics that the one period model abstracts from.

The sequential monetary friction is encoded by two inequalities, together with physical feasibility:

$$\begin{aligned} \text{Payment at } t=0: \quad & p\cdot\tau^{h,+} + q\beta^h \le m^h - q\ell^h && \text{(T)} \\ \text{Settlement at } t=1: \quad & p\cdot\tau^{h,-} + (b^h+\beta^h) + \ell^h + \mu^h \ge 0 && \text{(R)} \\ \text{Feasibility:} \quad & x^h \in \mathbb{R}^L_+, \quad \ell^h \le 0, \quad \beta^h \ge -b^h && \text{(F)} \end{aligned}$$

where $\mu^h \equiv m^h - q\ell^h - q\beta^h - p\cdot\tau^{h,+} \ge 0$ denotes money left unspent at $t=0$.

Constraint (T) is the payment constraint: purchases of goods and bonds at $t=0$ must be paid with money on hand—the endowment m^h plus the net resources raised from borrowing and bond sales.[10] Its right-hand side lists, in collapsed form, the instruments admissible in trade: in the baseline, money alone. Extending this list is precisely how Section 3 modifies the model. Sales proceeds cannot be recycled into purchases within $t=0$: in the model's timing, they arrive only at settlement.

Constraint (R) is the settlement condition, and it is where the monetary hierarchy resides: every financial obligation is discharged in outside money at $t=1$. The household's settlement resources are its sales proceeds $p\cdot\tau^{h,-}$, the delivery on its post-trade bond holdings $b^h+\beta^h$, and its unspent balances μ^h; against these stands the repayment $-\ell^h$ of its borrowing, at face value.

Constraint (F) collects the remaining restrictions: consumption is non-negative; households can borrow from the banking sector but cannot lend to it—lending occurs instead through bond purchases, $\beta^h > 0$; and bond sales are bounded by bond holdings.

In Dubey and Geanakoplos (2026) the carried-over money balances play a role because of the richer instrument set (the published lecture also analyses credit cards) and the finer time line (five intervals within the period).[11] In our case, feasible consumption plans are unaffected by whether unspent balances are carried or not, so we may drop μ^h from (R) and work with the tight settlement constraint

$$p\cdot\tau^{h,-} + (b^h+\beta^h) + \ell^h \ge 0. \tag{R$'$}$$

[10]The sign convention parallels that for loans: a bond sale is a negative net purchase, $\beta^h < 0$, so that $-q\beta^h > 0$ is money raised at $t=0$ against a delivery forgone at $t=1$, just as $\ell^h < 0$ is money $-q\ell^h > 0$ borrowed against a repayment $-\ell^h$ at $t=1$.

[11]In their time line, goods trade in the second interval, credit-card debts are settled in money in the third, and bank loans are repaid in the fourth. Money held back from the goods market can therefore be put to an intermediate use—paying a credit-card bill—and sales proceeds can be held back from the credit-card settlement to repay the bank. Whether agents do so is a genuine decision margin, and it ties the cash price and the credit-card price of each good together: Dubey and Geanakoplos show that in a robust equilibrium no money is hoarded across either interval, so that all of it is spent on the goods markets. In our two-date model there is no intermediate use: money not spent at $t=0$ can only meet settlement at $t=1$, where borrowing less does the same job at lower cost.

Lemma 1 in the appendix shows that this omission is without loss of generality. The lemma is a statement about financing, not about spending: both constraints depend on the household's financial choices only through the net position $\ell^h + \beta^h$, and idle cash held alongside outstanding bank debt is redundant bookkeeping—the household could reduce its borrowing until either the idle balance or the debt is exhausted, financing the same consumption bundle while easing settlement. If the household owes nothing, the tight constraint holds trivially, since every term on its left-hand side is non-negative. No consumption opportunity is gained or lost either way; the same bundle is merely financed more tightly.

Two observations complete the description. First, loans and bond trades convert future into current money at the same price q, so only the sum $\ell^h + \beta^h$ enters the household's constraints; the decomposition between the two is irrelevant to the household, though not to the aggregates, where bank credit and the bond market clear separately. Second, the bond market at $t = 0$ redistributes existing money between households—the buyer's payment capacity falls by exactly what the seller's rises—and creates none: aggregate payment capacity for goods is unaffected by bond trade. Only a change in the set of instruments admissible in (T) can expand what can pay. This distinction carries the analysis of Section 3.

The sequential structure—payment at $t = 0$, settlement at $t = 1$—is the decisive distinction from a Walrasian model. The impossibility of simultaneous exchange—the fact that a seller cannot immediately recycle the proceeds of a sale into a new purchase—is not a market imperfection but a physical constraint on economic activity: delivering a good and receiving a countervalue are acts that necessarily occur in sequence, not simultaneously. It is this irreducible asynchronicity of exchange that makes the existence of both a settlement medium and financial contracts a natural institutional response. The model provides a precise formal representation of our actual institutions of monetary exchange.

2.4. Budget set and household decisions

Given prices $(p, q) \in \mathbb{R}^L_{++} \times \mathbb{R}_+$, endowment $e^h \in \mathbb{R}^L_+$, outside-money holdings $m^h \in \mathbb{R}_+$ and bond holdings $b^h \geq 0$, define household h's *budget set* as

$$\mathcal{B}^h(p, q) \equiv \left\{(x, \ell, \beta) \in \mathbb{R}^L_+ \times \mathbb{R}_- \times [-b^h, \infty) \;\middle|\; (x, \ell, \beta) \text{ satisfy (T) and (R}'\text{)}\right\}. \quad (1)$$

The sign restrictions of (F) are carried by the domain: consumption is non-negative, $\ell \leq 0$, and bond sales are bounded by bond holdings.

Each household chooses (x^h, ℓ^h, β^h) to maximise utility subject to the constraints:

$$(x^h, \ell^h, \beta^h) \in \arg\max \left\{u^h(x) \;\middle|\; (x, \ell, \beta) \in \mathcal{B}^h(p, q)\right\}. \quad (2)$$

2.5. Collapsed representation

The two constraints defining the budget set can be collapsed into one. For any consumption bundle x^h, the financial position that maximises payment capacity at $t = 0$ is the largest borrowing compatible with settlement, $\psi^h = -(p \cdot \tau^{h,-} + b^h)$: the household borrows against everything it will have at settlement—its sales proceeds and its bond deliveries. Substituting this position into (T) yields the *collapsed representation*

$$p \cdot \tau^{h,+} - q\, p \cdot \tau^{h,-} \leq m^h + q\, b^h. \tag{3}$$

The collapse is exact: a consumption bundle is attainable in $\mathcal{B}^h(p, q)$ if and only if it satisfies (3). If the inequality holds, the maximal position supports the bundle; conversely, every element of the budget set has $\psi^h \geq -(p \cdot \tau^{h,-} + b^h)$, so (T) implies the inequality. The household's problem therefore reduces to maximising u^h subject to the single constraint (3).

This reveals a wedge between purchase and sale prices—Dubey and Geanakoplos' "wedge between buying and selling," the central friction of their analysis: purchases must be funded at unit price, while sales are effectively discounted by q because their proceeds arrive at $t = 1$ and can only be converted into $t = 0$ liquidity through borrowing. The bond endowment enters at qb^h, not b^h: a bond is money at a discount. The intercept $m^h + qb^h$ is the household's *effective liquid wealth*—cash plus bonds at their monetised value.

Relative to the bond-free economy, the bond shifts the budget frontier outward without altering its shape: bonds change the household's liquidity, not the friction it trades under.

The collapsed representation is more than a computational convenience; it exposes a structural feature of the theory. The model's premises are decidedly non-Walrasian: a nominal and a real sphere, exchange that cannot be synchronised, and trade against money organised market by market in the cash-in-advance tradition of Clower (1967). The projection onto consumption bundles returns a problem that is formally Walrasian again—a single budget constraint—except that prices are *directional*: a good costs p_k when bought and yields only qp_k when sold. The entire monetary friction is absorbed into this one asymmetry, a bid–ask spread proportional to the interest rate—a structure familiar from the transaction-cost equilibria of Foley (1970), where buying and selling prices likewise diverge, there because marketing consumes resources, here because settlement takes time.[12] This is what makes the model analysable despite its non-Walrasian foundations: the tools

[12] Since a round trip in the same good loses $(1 - q)p_k$ with certainty, a household trades a pair of goods only where its marginal gains from trade exceed the wedge. Executed trades therefore remain individually beneficial, while marginal trades are extinguished—the wedge operates like a tax on transactions. This is the household-level counterpart of the gains-to-trade hypothesis of Dubey and Geanakoplos (2026), who measure gains to trade net of precisely such a shrinkage of sales receipts, and who speak equivalently of an effective interest rate.

for handling Walrasian budget sets—demand theory, geometry, computation—apply once prices are read directionally. The numerical analysis at the end of this section exploits exactly this structure.

The parameter q governs the severity of the intertemporal friction. As $q \to 1$ (equivalently $r \to 0$), the wedge between purchase and sale prices closes and the budget set converges to a standard Walrasian budget set. Equilibrium allocations converge to Walrasian allocations and relative prices $p_k/p_{k'}$ converge to Walrasian exchange ratios. Yet individual money prices p_k diverge to infinity: the nominal scale is lost precisely as the monetary friction vanishes. This is the Walrasian indeterminacy of the price level reasserting itself at the boundary. It is outside money that pins down the price level for all $r > 0$; at $r = 0$, money has lost all scarcity and with it the power to anchor nominal values. Note that as $q \to 1$ the bond discount vanishes too—money and bonds become indistinguishable, so the composition of outside wealth stops mattering exactly when its level stops anchoring prices.

Two slacks of the constraint system provide the bridge from feasibility to the equilibrium accounting that follows. Writing $\psi^h \equiv \ell^h + \beta^h$ for the household's net financial position, define the slack variables

$$\begin{aligned} \sigma_0^h &\equiv m^h - q\psi^h - p \cdot \tau^{h,+} \;=\; \mu^h \;\geq\; 0, \\ \sigma_1^h &\equiv p \cdot \tau^{h,-} + b^h + \psi^h \;\geq\; 0. \end{aligned} \tag{4}$$

The variables σ_0^h and σ_1^h are the slacks of (T) and (R′), recording unspent liquidity at $t = 0$ and surplus settlement capacity at $t = 1$: a plan is feasible if and only if $(\sigma_0^h, \sigma_1^h) \in \mathbb{R}_+^2$; at an optimum both vanish, and summing the two zero-slack conditions across households delivers the central accounting identity of the model in the settlement-absorption subsection below. In this notation, Lemma 1 says that no consumption possibility is lost by driving the first slack to zero—a statement about the budget set, in which preferences play no role. That σ_0^h actually *is* zero at an optimum is the separate, preference-based argument made there.

The collapsed representation also yields a first aggregate statement, anticipating the equilibrium analysis. At an optimum the collapsed constraint (3) binds: were it slack, the household could purchase more of some good and, with strictly increasing preferences, be better off. Summing the binding constraints over households and using that aggregate purchases equal aggregate sales ($\sum_h p \cdot \tau^{h,+} = \sum_h p \cdot \tau^{h,-} \equiv V$, the nominal volume of trade) gives

$$(1 - q)\, V = M + qB : \tag{5}$$

the wedge on the volume of trade absorbs the economy's effective liquid wealth. This is the household-side face of the settlement absorption identity derived in the equilibrium analysis below.

2.6. Monetary equilibrium

A *monetary equilibrium* consists of prices $(\bar{p}, \bar{q}) \in \mathbb{R}^L_{++} \times (0,1)$ and an allocation $\{(\bar{x}^h, \bar{\ell}^h, \bar{\beta}^h)\}_{h\in\mathcal{H}}$ such that

(i) **Household optimality:** $(\bar{x}^h, \bar{\ell}^h, \bar{\beta}^h)$ solves (2) for all $h \in \mathcal{H}$,

(ii) **Goods-market clearing:** $\sum_{h\in\mathcal{H}}(\bar{x}^h - e^h) = 0$,

(iii) **Bond-market clearing:** $\sum_{h\in\mathcal{H}} \bar{\beta}^h = 0$,

(iv) **Credit market:** The banking sector accommodates aggregate borrowing $\bar{D} \equiv -\sum_{h\in\mathcal{H}} \bar{\ell}^h$ at the policy rate $\bar{q}$.

Note that the decomposition of $\bar{\psi}^h$ into $(\bar{\ell}^h, \bar{\beta}^h)$ is not unique—households are indifferent between borrowing and bond sales at the common price $\bar{q}$—but all equilibrium objects that matter are invariant to the decomposition: the consumption allocation, the net positions $\bar{\psi}^h$, prices, and, by bond-market clearing, aggregate credit $\bar{D} = -\sum_h \bar{\psi}^h$.

The credit price as a policy instrument. In this model, q summarises the terms on which the banking system converts future settlement obligations into current payment capacity—ultimately anchored by the monetary policy stance. A decrease in q (equivalently, an increase in the interest rate $r = (1-q)/q$) corresponds to monetary tightening: borrowing becomes more expensive and the wedge between purchase and sale prices widens. An increase in q (a decrease in r) corresponds to easing: the cost of bridging the asynchronicity of exchange falls. Throughout, q is the policy instrument and $\bar{D}$ emerges in equilibrium, as discussed in the remark on instruments above; studying shifts in q is a reduced-form way to discuss monetary policy.

2.7. Settlement absorption

At an optimum, both slacks of the constraint system vanish. That unspent liquidity σ_0^h is zero follows as in the collapsed representation: money left idle could buy more of some good. That surplus settlement capacity σ_1^h is zero follows by the mirror argument: settlement resources beyond obligations are worthless at the horizon, so the household converts them into purchases at $t = 0$ by lowering ψ^h. Summing the two zero-slack conditions across households and using goods- and bond-market clearing yields two aggregate statements: $V = M + q\bar{D}$—all money on hand is spent on goods—and $\bar{D} = V + B$—all settlement resources, sales proceeds and bond deliveries alike, are pledged to the banking sector. Combining the two gives the *settlement absorption condition*:

$$(1-q)\,\bar{D} = M + B. \tag{6}$$

Substituting $\bar{D} = V + B$ back recovers the volume identity (5): the two are the same accounting, read once from the credit side and once from the trade side.

The condition admits several readings. Writing $r = (1 - q)/q$ for the nominal interest rate, it reads $r \cdot (q\bar{D}) = M + B$: the interest on inside money absorbs the economy's *entire* outside wealth, money and bonds alike. This is how the framework resolves the puzzle of why intrinsically valueless money is accepted in a finite horizon—the Hahn problem (Dubey and Geanakoplos, 2026, §3.5). Borrowers owe the banking sector more than they borrowed; to obtain the difference they sell goods to the holders of endowed money and bonds, and at settlement the entire outside wealth is delivered to the public sector as interest. The question of who holds the money at the end has a definite answer: the consolidated public sector.

At the individual level, outside money and inside money are functionally equivalent for payment. At the aggregate level, they play fundamentally different roles: inside money facilitates payment but creates obligations; outside money is the ultimate settlement medium. The identity (6) captures this asymmetry: every unit of inside money must ultimately be retired with outside money—recycled sales proceeds and bond deliveries.

The absorption condition is the central equilibrium condition of Dubey and Geanakoplos (2026) in this paper's notation. Because our notation departs from theirs at several points, the correspondence is collected in Table 1.

Object	This paper	Dubey–Geanakoplos
Outside money endowment	M	$\bar{m}$
Bond endowment	B	$\bar{b}$
Bank money (loans extended)	$q\bar{D}$	M
Nominal interest rate	$r = (1 - q)/q$	r
Outside wealth	$W = M + B$	fiat wealth $\bar{m} + \bar{b}$
Absorption condition	$(1 - q)\bar{D} = M + B$	$rM = \bar{m} + \bar{b}$

Table 1: Notation correspondence with Dubey and Geanakoplos (2026). Note the inversions: their M is bank money, our M is the outside money endowment.

Aggregate payment capacity for goods at $t = 0$ equals $M + q\bar{D}$: under bond-market clearing, bond purchases and sales net out, so the bond stock contributes to payment capacity only through its monetised value inside $\bar{D}$. Nominal prices adjust so that desired real trades can be financed with available liquidity. Any monetary innovation that expands aggregate payment capacity is inflationary; one that merely relabels existing payment instruments is not. It is precisely this baseline neutrality of the bond stock that the stablecoin layer of Section 3 will break.

2.8. Existence and gains from trade

Dubey and Geanakoplos (1992) and Dubey and Geanakoplos (2026) show that, under the stated assumptions, a monetary equilibrium exists if and only if a gains-from-trade condition is satisfied. The key quantity is $\gamma(e)$, which measures the economy's gains from trade: how much the terms of exchange can be worsened before households no longer find it worthwhile to trade.[13] The existence theorem states: monetary equilibrium exists if and only if

$$\gamma(e) > r = \frac{1-q}{q}.$$

Since our economy is an instance of theirs—the parameter settings and the relabelling of the timeline are collected in Remark 1 in the appendix—the theorem applies directly. With bank credit unrestricted, their ownership condition (every good owned by some agent with unbounded credit) reduces to our assumption $\sum_h e^h \gg 0$.

The interest rate creates a wedge between buying and selling prices; if the economy's gains from trade are insufficient to overcome this wedge, no household wishes to borrow, money has no value, and the economy collapses to autarky. The economically relevant region thus lies between two boundaries: the Walrasian limit $q \to 1$, where money loses its scarcity and the price level its nominal anchor, and the gains-from-trade threshold, below which monetary trade collapses altogether.

The existence threshold and the absorption condition are two readings of the same fact: by (6), $r = (M + B)/(q\bar{D})$ —the existence threshold of the published version, $\gamma^* = (\bar{m} + \bar{b})/M$, in our notation (cf. Table 1). The economy must generate enough gains from trade to carry the interest burden that, in equilibrium, transfers its entire outside wealth to the public sector.

2.9. Interpreting the model: geometry and computation

The mechanics of an abstract model are absorbed in different ways by different readers: some think in pictures, others by computing examples. This subsection offers both readings. The two figures are not interchangeable illustrations; they depict the two spheres from which the economy is built. The net-trade diagram (Figure 1) lives in the real sphere; it is kin to the Edgeworth-box analysis of the published lecture (Dubey and Geanakoplos, 2026, Figure 3), recentred at the endowment and drawn for a single household. The transfer diagram (Figure 2) lives in the nominal sphere, where payment and settlement reside; it

[13] Formally, $\gamma(e)$ is defined as the supremum of all $\gamma \geq 0$ such that there exist feasible net trades $\{\tau^h\}_{h\in\mathcal{H}}$ with $\sum_h \tau^h = 0$ and $u^h(e^h + \tau^h) > u^h(e^h)$ for all h, even when each household's receivables $p \cdot \tau^{h,-}$ are shrunk by the factor $1/(1+\gamma)$. The shrinkage penalises the sale side of trade, mimicking the discount imposed by the interest rate.

is this paper's device, adapted from the geometric analysis of general equilibrium with incomplete markets (Magill and Quinzii, 1996). What connects the two spheres is the price system: for given (p, q), every trading plan in the real sphere has an image in the nominal sphere, and a plan is feasible precisely when its image lies in the feasibility set $\mathcal{F}^h_p$. The collapsed representation introduced earlier is the statement that, for the simple settlement structure studied here, the nominal sphere can be folded back into the real one—compressed into directional prices—so that equilibrium can be found in the real sphere alone.[14]

Geometry. The collapsed constraint (3) defines a characteristic kinked budget set in net-trade space. Figure 1 illustrates for $L = 2$. For any pair of goods (k, k'), the household faces a different effective price ratio depending on whether a good is purchased or sold. A good that is purchased costs its full price p_k; a good that is sold contributes only qp_k to the household's purchasing power. This creates kinks at the regime boundaries—the points where a household switches from being a net buyer to a net seller of a particular good. In general, the budget frontier consists of distinct linear facets, one for each pattern of net purchases and net sales across goods, plus the boundary where the household does not borrow at all.

[14] Whether richer settlement structures—several settlement assets, tiered finality in the sense of Mehrling (2013)—admit such a projection is not obvious, and is a question for separate work.

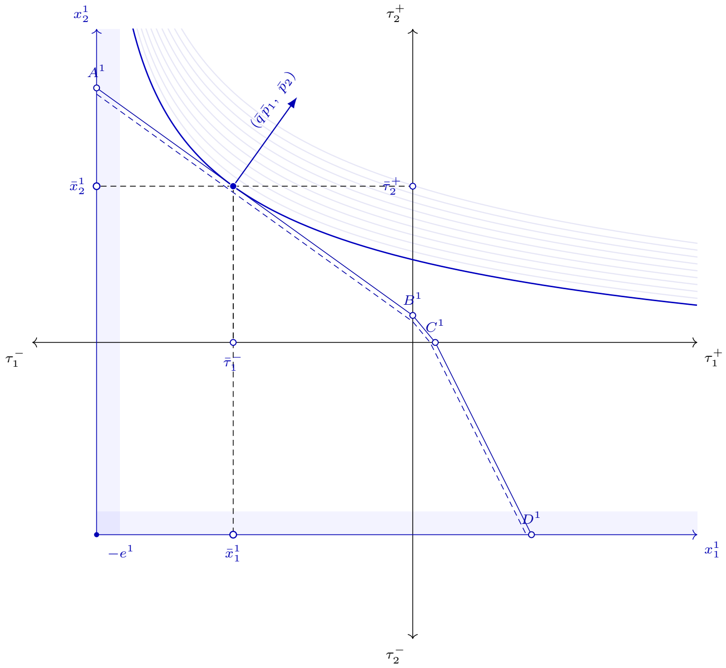

Figure 1: Budget-feasible net trades for household h under the collapsed budget constraint $\bar{p}\cdot\tau^{h,+} - \bar{q}\,\bar{p}\cdot\tau^{h,-} \leq m^h + \bar{q}\,b^h$ ($L = 2$), illustrated for a household with $b^h > 0$. The solid frontier is the budget frontier at effective liquid wealth $m^h + \bar{q}\,b^h$; the dashed frontier is the same household's frontier without bonds: the bond endowment shifts the frontier outward by $\bar{q}\,b^h$ without altering its shape. The kink points A, B, C, D mark boundaries between trading regimes. On segment A–B, the household is a net seller of good 1 and net buyer of good 2: sales revenue from good 1 is discounted by $\bar{q}$ because proceeds arrive at settlement ($t = 1$) and can only be converted to purchasing power at $t = 0$ through borrowing. On segment B–C, both goods are purchased using only outside money and credit; the effective price ratio equals the undiscounted ratio p_1/p_2. On segment C–D, good 2 is sold and good 1 purchased; sales revenue from good 2 is discounted by $\bar{q}$. The optimum $\bar{\tau}^h$ lies on segment A–B.

The constraints admit an equally transparent picture in *net nominal income transfer space*, where the transfer diagram literally plots the constraint system. For a given net trade τ^h, the financial position moves the household along a ray of slope $-1/q$ in (ν_0, ν_1)-space: each additional unit of borrowing (ψ^h decreasing by one) adds q to ν_0^h and subtracts one from ν_1^h—the financial system converts settlement capacity into payment capacity at rate q. The feasible set of net nominal income transfers, $\mathcal{F}_p^h(m^h, b^h, q) \subseteq \mathbb{R}^2$, is the shaded region in Figure 2. Idle money is the horizontal distance from the $\nu_0 = 0$ axis; that no consumption possibility is lost by eliminating it is the geometric content of Lemma 1.

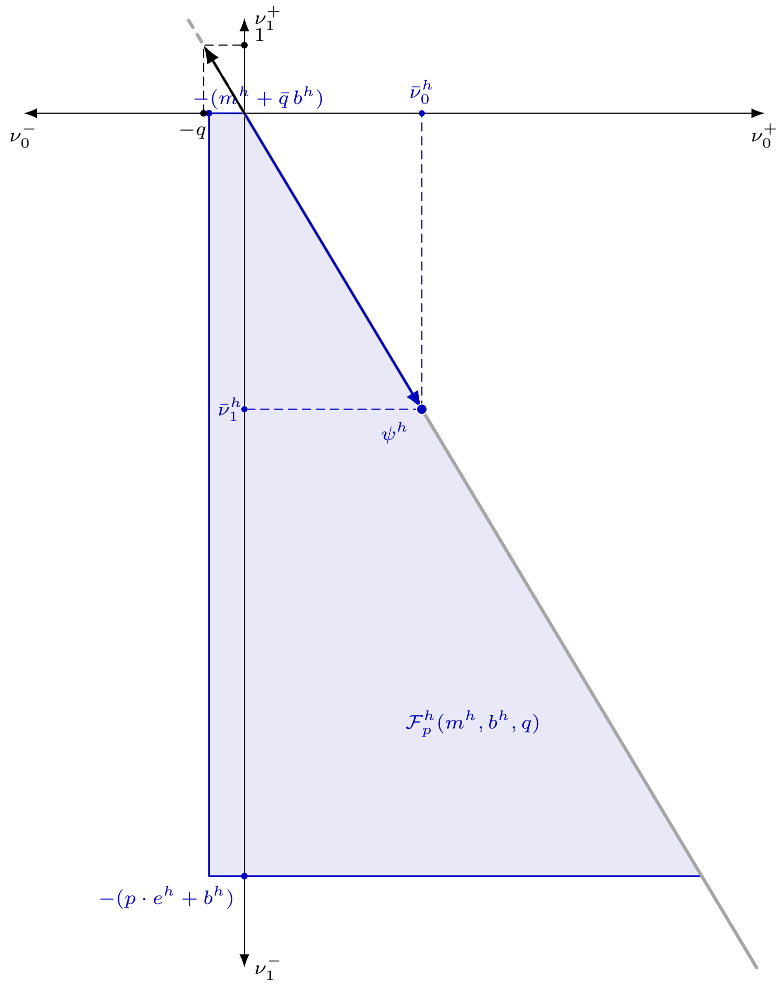

Figure 2: Household h's feasibility set $\mathcal{F}_p^h(m^h, b^h, q)$ in net nominal income transfer space, illustrated for a household with $b^h > 0$. The axes show the net nominal income transfers (ν_0, ν_1) at the two dates as signed quantities, so the feasible set (shaded) is not confined to the non-negative quadrant: $\mathcal{F}_p^h(m^h, b^h, q) \subseteq \mathbb{R}^2$. The boundary marks at $-(m^h + \bar{q}\, b^h)$ (horizontal axis) and $-(p \cdot e^h + b^h)$ (vertical axis) are the maximal deficits of the household: with no sales and no credit it can spend at most its cash plus the monetised value of its bonds, and selling its entire endowment it commands at most its sales proceeds plus its bond delivery in settlement capacity. For $b^h = 0$ the marks reduce to $-m^h$ and $-p \cdot e^h$. The financial position converts settlement capacity into payment capacity at rate q (ray slope $-1/q$).

Computation. The computational reading is provided by a companion R package, `dgme`, which implements the economy of this section for two-household, two-good Cobb–Douglas examples.[15] The solver puts the collapsed representation to work: it evaluates the direction-

[15] Openly archived at `https://github.com/Martin-Summer-1090/dgme`; the release tagged v0.5.0 corresponds to this version of the paper. The package computes equilibria, verifies the identities of this section numerically, and generates the figures shown here; its vignettes reproduce all numerical examples of this paper and of the published Klein Lecture, and contain further stablecoin exhibits not included here for reasons of length: sweeps over issuance scale and reserve composition, calibrated issuer scenarios, and the price-level continuum of the two-money economy of Section 4.2. The reader is invited to experiment.

dependent demands facet by facet and solves the resulting Walrasian-style excess-demand system for money prices—no further monetary structure is needed, which is the projection argument above in practical form. The bond accounting is reconstructed afterwards: positions ψ^h, trade volume V, credit $\bar{D} = V + B$. In the computed equilibria the identities (5) and (6) and the decomposition $\bar{D} = V + B$ hold to 10^{-6}, and the equivalence between the bond economy and the bond-free economy with money endowments $m^h + qb^h$ holds to machine precision: prices and allocations coincide exactly; only the accounting differs.

The examples recover, in miniature, the monetary-policy discussion of the published lecture. In the parameterisation behind the figures ($\bar{q} = 0.6$, so $r = 2/3$), endowing households with bonds worth half the money stock ($B = M/2$) raises all money prices by exactly the factor $(M + \bar{q}B)/M = 1.3$ while leaving the allocation unchanged: effective liquid wealth sets the nominal scale, and its composition does not matter—the baseline neutrality that Section 3 will put under pressure. Easing policy ($q \to 1$) raises nominal prices and credit while utilities rise toward their Walrasian levels; tightening widens the wedge until, at the gains-from-trade threshold, trade collapses. None of this is new relative to Dubey and Geanakoplos (2026); the point is that the reader can generate it, vary it, and check every identity of this section along the way.

3. How stablecoins can raise the price level

The Dubey and Geanakoplos (2026) model developed in Section 2, now gives us the language to express the channels through which stablecoins can raise the price level and thereby exert inflationary pressure. This is the leading question of this section.[16] The key quantity is the economy's effective payment capacity, which we write as

$$M + qD + \theta B. \tag{7}$$

The parameter $\theta \in [0, 1]$ is new here: it expresses the idea that part of the stock of government bonds B can be transformed into current payment capacity at $t = 0$. Let us call it the moneyness factor of government bonds. In the baseline of Section 2, $\theta = 0$: a bond contributes to liquidity only indirectly, in that a household can borrow against

[16]The framework is essentially static: a one-period economy in which prices clear markets at each configuration of parameters. It cannot generate inflation in the dynamic sense of a sustained rate of price increase. When we say that stablecoins are "inflationary," we mean precisely that the equilibrium price level is higher when the stablecoin layer is present than when it is absent—a comparative statics statement about the *level* of prices, not their *rate of change*. This usage follows Dubey and Geanakoplos (2026). The comparative static identifies genuine economic mechanisms: payment innovations that expand aggregate purchasing power exert upward pressure on prices. Whether this manifests as a one-time level shift or as sustained price growth depends on dynamics outside the model, but the direction of the effect is unambiguous.

the bond's delivery at $t = 1$, converting it into qb^h units of current liquidity through the banking sector's $(-q, 1)$ loan technology. In the aggregate this contribution is therefore part of the credit term qD, not a separate term.

We argue that stablecoins are a form of financial engineering which implements this moneyness of government bonds, thereby adding to current payment capacity. Once in place it not only mobilises a stock of deferred claims to outside money into liquidity at $t = 0$ – the stock channel – it also amplifies inflationary effects of newly issued public debt – the flow channel.

Each of the channels studied below is a movement of the moneyness factor θ in (7) or movements in the bond stock B. What the accounting contributes is not a separate term for each channel but a way of telling them apart: it separates what the stablecoin adds from what was there before. A larger bond stock raises payment capacity even at $\theta = 0$, through the monetised value of the additional bonds inside qD—a fiscal effect the baseline economy already contains, before any stablecoin exists. The stablecoin changes nothing but the bonds' moneyness. Read this way, the analysis delivers a taxonomy of channels within a single accounting, not a bundle of separate issues—and the accounting also keeps visible what does *not* move: the settlement obligation that delivers the economy's outside wealth to the public sector, as Corollary 1 will show.

The section proceeds along the margins of θB. Section 3.1 states the moneyness margin in its purest form: Dubey and Geanakoplos' own theorem that prices rise when bonds become money, translated into the two-date framework and given its institutional content—the stablecoin enters the paper here, as the piece of financial engineering that makes the theorem's hypothesis real. Section 3.2 then endogenises the parameter through the institution that carries it: the issuer's balance sheet, with its scale S and its reserve composition ρ. Section 3.3 lets the bond stock itself respond—the fiscal margin—and shows that the resulting channel is not a second mechanism but an amplification of a comparative static the baseline already contains. The breadth of adoption—how much of goods trade the coin can pay for—completes the taxonomy at the end of the section. Two neighbouring questions are deliberately deferred. How monetary policy can respond to the channels is a response, not a channel, and has its own place (Section 5). And a coin backed by assets that are nobody's liability is not a movement of θB at all, but a candidate competitor to the monetary arrangement itself; that escalation is the subject of Section 4.

Two remarks on scope. First, the analysis studies stablecoins at the margin where they pay for goods and services in general—the margin the payment constraint of the model formalises. Reaching that margin takes more than the coin itself: payments are a two-sided market. Financial engineering can supply an instrument fit to pay, but an instrument pays for goods only where sellers are equipped and willing to receive it—acceptance at the checkout, the merchant's processing infrastructure, the plumbing that connects a token to an invoice. The model compresses this second leg into a single standing assumption: by admitting the coin into the payment constraint (T), it takes the acceptance infrastructure for goods purchases as already in place. That is today largely counterfactual: existing coins circulate overwhelmingly within crypto-asset markets, and the scenario in which they buy ordinary goods is a possibility contemplated by their design and by the regulation now being written for them, not a description of current circulation. It is precisely this possibility the section analyses; every channel derived below is conditional on the acceptance leg being built. Whether and how easily it can be built is not a natural constant but itself a policy margin—regulation and infrastructure policy decide whether coins stay confined to the crypto universe or reach the checkout—and it returns as such among the policy postures of Section 6. Second, run risk, redemption dynamics, and the composition of bank funding are financial-stability questions with a literature of their own; the paper abstracts from them throughout.

3.1. The stock channel: When bonds become money

The stock channel is Dubey and Geanakoplos' own result. Theorem 4 of the published lecture carries its content in its name: "When Bonds Become Money, Prices Rise" (Dubey and Geanakoplos, 2026, §6.4). The theorem is proved in an economy richer than our baseline—it contains credit cards as a second deferred payment instrument—so we restate it informally and in the vocabulary of this paper, and translate it into the two-date framework below; for the exact statement see Dubey and Geanakoplos (2026, §6.4).

Theorem (Dubey–Geanakoplos, Theorem 4; informal restatement)**.** *Consider a monetary economy in which households are endowed with government bonds but bonds cannot be used to pay for goods, and suppose it has a monetary equilibrium. If bonds become directly usable for purchases—the economy otherwise unchanged—then the new economy has a monetary equilibrium in which every household consumes exactly the same bundle as before, and the money price of every good is higher. Goods bought against bonds trade at the deferred-settlement price: the money price marked up by the interest-rate wedge.*[17]

Let us translate this result into our simplified two-date model, without credit cards and with government bonds as the only additional instrument. One point should be made precise first, because the theorem's name can obscure it: "bonds become money" does not mean that bonds move from illiquid to liquid. In the theorem's economy just as in

[17]The last clause is where the credit cards of their economy appear in the original: Dubey and Geanakoplos state it as "the bond price of each commodity is equal to its credit card price." A credit card charge and a bond are both promises of money at settlement, so both buy goods at the same deferred price—the money price marked up by the effective interest-rate wedge. Our baseline switches credit cards off (Remark 1); the deferred-settlement reading given in the text is what survives the specialization.

our baseline, households can borrow against a bond's delivery, so a bond that cannot pay for goods is nevertheless liquid at discount q through the banking sector's $(-q, 1)$ technology, its contribution to payment capacity recorded inside qD in (7). What the theorem's hypothesis changes is accordingly more specific than the name suggests: it moves the marginal bond from bank-intermediated liquidity at q to direct spendability at face value—a gain of $1 - q$ per unit of face value, and precisely the movement of the term θB in (7).

The question the theorem raises is then institutional: what arrangement would make bonds usable as money? Bonds, as they currently exist, are not payment instruments. They sit in custody accounts and move through tiered securities-settlement systems rather than from hand to hand; they come in denominations unsuited to everyday transactions; households, as the endowment discussion of Section 2.2 noted, mostly do not even hold them directly, but through fund shares that must first be redeemed into bank money; their market price fluctuates with the interest rate; and accrued interest makes even the unit in which they are quoted shift from day to day. None of these frictions exists inside the model—which is why the theorem can make bonds money by assumption. Implementing moneyness of government bonds in reality requires financial engineering.

A stablecoin backed by short-term Treasuries is exactly this piece of financial engineering, and it removes the frictions one by one: tokenisation makes the claim divisible and transferable peer to peer; the peg to the currency unit provides a stable unit for quoting prices; redemption on demand pins the peg. Each feature has a model counterpart. Tokenisation opens the market pair bond–goods that the baseline keeps closed (Remark 1); the peg means the coin enters the payment constraint (T) at face value; and acceptance at par by sellers is consistent with optimising behaviour because *all* sales proceeds—money and coin alike—become available only at settlement, where the coin delivers exactly one unit of outside money. Dubey and Geanakoplos note themselves that in a one-period model the distinction between purchases with bond-backed stablecoins and direct purchases with government bonds is insignificant (their §2.7): the stablecoin is the institutional arrangement that makes the hypothesis of their theorem real.

One caution belongs here, recalling the scope remark that opened the section: the engineering supplies the instrument, not the market. The theorem's hypothesis becomes real only where the acceptance leg stands as well.[18]

Since the lecture discusses stablecoins only in passing, and mostly through the lens of its credit-card apparatus, it is worth stating precisely how the two instruments relate. The assimilation is accurate at exactly one level: a credit-card charge and a bond-backed

[18] The second leg is a policy margin in its own right. Public initiatives that build acceptance infrastructure for coin payments—or regulation that eases its build-out—switch on the very channel this section derives; among the three postures of Section 6 they belong to *amplify*, with the price-level consequences that follow.

coin are both promises of settlement money at $t = 1$, so both buy goods at the same deferred-settlement price—the fact Theorem 4 turns on. Beneath that level, the two expand payment capacity through different mechanisms. A credit card defers the payment for a purchase to settlement, where sales proceeds can meet it; the same money thereby finances more trade within the period—an increase in the velocity of the existing payment media. A bond-backed coin creates no deferral: it takes a deferred claim that already exists—the bond, outside money payable tomorrow—and makes it liquid for payments already today. In the accounting of (7), the credit card raises the trade a given $M + qD$ can finance, while the coin adds the new term θB. This is why our specialization can switch credit cards off (Remark 1) and still describe the stock channel fully.

Formally, the stablecoin layer changes the constraint system of Section 2 in one place. Let $s^h \in [0, \theta b^h]$ denote the face value of bonds that become spendable in payments at $t = 0$ for household h by the presence of stablecoins. The payment constraint now admits the stablecoin alongside money,

$$p \cdot \tau^{h,+} + q\beta^h \;\leq\; m^h - q\ell^h + s^h, \tag{T_θ}$$

the settlement constraint records that bonds backing the stablecoins deliver to their new holders,

$$p \cdot \tau^{h,-} + (b^h - s^h + \beta^h) + \ell^h \;\geq\; 0, \tag{R'_θ}$$

and the no-short-sales bound tightens to $\beta^h \geq -(b^h - s^h)$: a bond tendered as a coin cannot also be sold. The bound $s^h \leq \theta b^h$ applies to the bond *endowment*, not to post-trade holdings: the coins outstanding are minted against the existing stock before trade begins. This is not a technicality. If bonds purchased at $t = 0$ could be coined within the trading date, buying a bond at q and spending it at par would be an arbitrage, and the price system would collapse. (Coin issuance against newly acquired bonds is the business of the issuer, and is taken up when θ is endogenised below.)

Two observations follow directly. First, for a household that tenders s^h, the constraints above are *exactly* the baseline constraints of a household endowed with $(m^h + s^h,\ b^h - s^h)$: tendering a coin converts a bond into money, at par, unit for unit. Second, since each unit tendered raises effective liquid wealth by $1 - q > 0$, households use the facility to capacity, $s^h = \theta b^h$—the model's rendering of Dubey and Geanakoplos' observation that agents strictly prefer to spend each Treasury rather than hold it as an investment (their §6.4). The parameter θ is therefore a binding primitive, not an upper bound. These two observations drive the following result, proved in Appendix A.

Proposition 1 (The stock channel)**.** *Fix $q \in (0, 1)$ and let $\theta \in [0, 1]$. Consider the economy $\mathcal{E}(\theta)$ in which each household h can use a fraction θ of its bond endowment b^h to pay for goods at $t = 0$ at face value—the baseline of Section 2 with (T) and (R′) replaced by* (T_θ) *and* (R'_θ) *and the bound $\beta^h \geq -(b^h - s^h)$.*

(i) (Equivalence.) *$\mathcal{E}(\theta)$ has the same monetary equilibria as the baseline economy with endowments $\big(m^h + \theta b^h,\ (1-\theta)b^h\big)_{h\in\mathcal{H}}$: making bonds spendable at par is equivalent to converting them into money, one for one, in the proportion θ.*

(ii) (Payment capacity and nominal volume.) *In any monetary equilibrium of $\mathcal{E}(\theta)$, the effective liquid wealth of household h is $m^h + qb^h + \theta(1-q)b^h$, and the nominal volume of trade satisfies*

$$(1-q)\,V \;=\; M + qB + \theta(1-q)B,$$

strictly increasing in θ whenever $B > 0$.

(iii) (Price level.) *Suppose portfolio composition is uniform across households: $b^h/(m^h + qb^h)$ is the same for all h. If the baseline ($\theta = 0$) has a monetary equilibrium with prices $\bar{p}$ and allocation $(\bar{x}^h)_h$, then $\mathcal{E}(\theta)$ has a monetary equilibrium with the* same *allocation and all money prices scaled by the factor*

$$\kappa(\theta) \;=\; 1 + \theta\,\frac{(1-q)B}{M + qB},$$

which exceeds one whenever $\theta > 0$ and $B > 0$.

Corollary 1 (Settlement absorption with stablecoins)**.** *In any monetary equilibrium of $\mathcal{E}(\theta)$, aggregate credit is $\bar{D} = V + (1-\theta)B$ and the settlement absorption condition is unchanged: $(1-q)\bar{D} = M + B$. Stablecoins change the composition of payments and the volume of trade, but not the settlement obligation that delivers the economy's outside wealth to the public sector.*

The proposition's parts differ in what they require, and the difference is economics, not bookkeeping. Parts (i) and (ii) are distribution-free: however the bond stock is spread across households, making it spendable raises the economy's payment capacity and its nominal volume of trade, strictly in θ. Part (iii) is not, and the reason can be put simply. To its holder, spendability is a windfall: a bond commanded q units of current purchasing power through the credit system and commands 1 once it pays at par—a gain of $1-q$ per bond, accruing to whoever holds bonds. Whether a windfall moves anything real depends on how it is spread. If it raises every household's liquid wealth in the same proportion—this is what the uniformity hypothesis of part (iii) says—it works like a change of units: all budgets expand together, prices rise by the common factor $\kappa(\theta)$, and no household's command over goods changes relative to any other's. If bond holdings are skewed, the same aggregate windfall lands unevenly: bond-rich households gain real purchasing power at the expense of money-rich ones, demand shifts toward the goods they prefer, and relative prices and the allocation move with it—Panel B of Table 2 below puts numbers to it.

This wealth effect is also the precise point at which the par-redeemable coin differs from the bond of Theorem 4 in Dubey and Geanakoplos (2026), and so marks what the proposition adds to the theorem it translates. In Dubey and Geanakoplos' theorem the allocation is unchanged *without any* hypothesis on portfolios, because their spendable bond generates no windfall in the first place: it buys goods at the deferred-settlement price—the money price marked up by the interest wedge—so spending a bond is worth exactly what monetising it through the credit system is worth, and no household gains from the new option. A par-redeemable stablecoin is more money-like than that: it hands its holder the full gain $1 - q$, and the wealth effect is the price of the extra moneyness—a consequence of the coin's design, not of the framework. The uniformity hypothesis of part (iii) is the deliberately special case that switches the effect off: when portfolio composition is uniform, the redistribution vanishes, and what remains is the pure nominal effect—prices scaled by the single factor $\kappa(\theta)$.

The proposition also identifies the channel in the empirical sense of the word: it says where to look for its footprint, and where not to. The price level in Proposition 1 responds to θ, the fraction of the bond stock that the payment constraint accepts at par—and to nothing else about the bonds' private holders. This is what separates the stablecoin from the money market fund, the institution that has intermediated essentially the same portfolio of short-term Treasuries for four decades and at far larger scale. The fund transforms T-bills into shares, and shares are not payment instruments: a holder who wants to buy goods must first redeem them into bank money, so every purchase still passes through the payment constraint (T) with money alone on its right-hand side. In the model's terms, θ stays at zero however large the fund sector grows. The stablecoin transforms the same T-bills into a token that the payment constraint (T_θ) accepts directly, and it is this piece of engineering—not the portfolio behind it—that moves θ off zero. Same assets, different engineering, different moneyness: the stock channel is payment-side, not portfolio-side. The rise of the money market funds is then not a counterexample to the mechanism but its natural control experiment. The fund sector expanded private T-bill intermediation to trillions of dollars without ever making the bills usable in payment—a movement along $\theta = 0$. For such a movement the proposition predicts exactly what was observed: no effect on the price level.

Proposition 1 also answers the second half of the section's question: what determines the size of the stock channel. The price-level effect is governed by a single ratio,

$$\theta \, \frac{(1-q)B}{M+qB},$$

and each of its determinants is influenced by a different actor. The parameter θ records how much of the bond stock the coin's design and its adoption make spendable—the regulatory and technological margin. The stock B is a fiscal magnitude: the channel needs bonds to monetise, and it is the fiscal authority that supplies them. The factor

$1-q$ is the monetary policy stance: the tighter policy, the wider the gap between par and bank-intermediated liquidity—and that gap is exactly what the coin monetises, so a given θ moves prices more in a tight-money regime. Under the equivalent quantity formulation of the remark on instruments in Section 2—the bank fixes the volume of new loans, and the interest rate emerges as the price that rations them—the same point reads: the coin adds payment capacity that never passes through the loan market, so spendable bonds bypass the rationing through which a tight policy binds. In the running example of Section 2 ($q = 0.6$, $B = M/2$), full spendability raises the price level by $2/13 \approx 15.4$ percent. The determinants also interact: when new debt is issued into an economy with an established stablecoin sector, the marginal bond arrives partly spendable, contributing $q + \theta(1-q)$ per unit to payment capacity at the moment of issue rather than the q of a bond that must be monetised through the credit system—deficits become more inflationary per bond issued as θ rises. We return to this interaction once θ is endogenised through issuance scale and reserve composition in the next subsection, where the issuer's balance sheet enters the model.

These statements can be checked computationally, and the check also quantifies the wealth effect just discussed. Table 2 recomputes the running example of Section 2 under the stablecoin layer for $\theta \in \{0, \frac{1}{2}, 1\}$. The computation exploits the equivalence of Proposition 1(i)—$\mathcal{E}(\theta)$ is solved as the baseline economy re-endowed with $(m^h + \theta b^h, (1-\theta)b^h)$—so what the numbers test are the identities of part (ii) and Corollary 1, and the scaling prediction of part (iii).

Panel A is the running example, whose bond distribution $b^h = m^h/2$ satisfies the uniformity hypothesis of part (iii). The prediction is borne out exactly: both money prices scale by the predicted factor $\kappa(\theta)$ to solver precision, the allocation does not move, and all three identities hold at the tolerance of the equilibrium computation. Two features of the aggregates deserve notice. Nominal trade volume rises with θ, from 7.15 to 8.25; aggregate credit does not move at all, $\bar{D} = (M+B)/(1-q) = 8.25$ throughout. The coin adds payment capacity and trade volume without adding settlement obligations—the computational face of Corollary 1. Panel B then redistributes the same bond stock towards the money-poor household ($b = (0.9, 0.2)$), so that portfolio composition is no longer uniform. The distribution-free statements are unaffected: all identities continue to hold exactly, and V and $\bar{D}$ take the same values as in Panel A. But prices no longer scale uniformly—at $\theta = 1$ the two goods' price factors, 1.149 and 1.173, straddle the formula's 1.154—and the allocation shifts by up to 0.039 units of consumption: the bond-rich household's gain tilts demand towards the good it favours. This is the redistribution discussed after Proposition 1 made quantitative: outside the uniform case, the stock channel moves relative prices and real consumption, not just the price level.

Table 2: The stock channel in the running example. Each row makes a fraction θ of the bond stock spendable at par, computed via the equivalence of Proposition 1(i) (endowments $m^h + \theta b^h$, $(1-\theta)b^h$; dgme package v0.4.0). Parameters: $q = 0.6$, $M = 2.2$, $B = 1.1$, so the predicted price factor is $\kappa(\theta) = 1 + \theta(1-q)B/(M+qB) = 1 + 2\theta/13$. p_1, p_2: money prices; $p_l(\theta)/p_l(0)$: observed price factors; V: nominal trade volume; $\bar{D}$: aggregate credit; max $|\Delta x|$: largest change in any household's consumption of any good relative to $\theta = 0$; residual: largest absolute violation of the three equilibrium identities $(1-q)V = M + qB + \theta(1-q)B$, $\bar{D} = V + (1-\theta)B$ and $(1-q)\bar{D} = M + B$. Panel A: the running example's bond distribution $b = (0.2, 0.9) = m/2$, so portfolio composition is uniform and Proposition 1(iii) applies. Panel B: identical aggregates but bonds redistributed, $b = (0.9, 0.2)$; composition is no longer uniform, so prices need not scale uniformly and the allocation shifts (the redistribution discussed after Proposition 1), while the identities of Proposition 1(ii) and Corollary 1 continue to hold exactly.

θ	p_1	p_2	$\frac{p_1(\theta)}{p_1(0)}$	$\frac{p_2(\theta)}{p_2(0)}$	$\kappa(\theta)$	V	$\bar{D}$	max $\|\Delta x\|$	residual
Panel A: the running example's bond distribution ($b^h = m^h/2$, uniform composition)									
0.00	2.606	2.176	1.0000	1.0000	1.0000	7.15	8.25	0.000	$4.5e-12$
0.50	2.806	2.344	1.0769	1.0769	1.0769	7.70	8.25	0.000	$5.1e-12$
1.00	3.007	2.511	1.1538	1.1538	1.1538	8.25	8.25	0.000	$5.3e-12$
Panel B: bonds redistributed ($b = (0.9, 0.2)$, non-uniform composition)									
0.00	2.580	2.260	1.0000	1.0000	1.0000	7.15	8.25	0.000	$< 10^{-12}$
0.50	2.772	2.456	1.0744	1.0865	1.0769	7.70	8.25	0.021	$< 10^{-12}$
1.00	2.964	2.651	1.1487	1.1729	1.1538	8.25	8.25	0.039	$< 10^{-12}$

Which of the two transformations, then, models the monetisation that actual stablecoins perform—the deferred-settlement bond of Dubey and Geanakoplos' theorem, or the par-spendable coin of the proposition? For stablecoins the question has a clear answer: the par case. A coin pegged at one dollar buys one dollar's worth of goods at the checkout; no merchant quotes a coin price marked up by the bill rate. The deferred-settlement price is the price of a promise whose holder is made to wait—a bond tendered directly in trade, or the credit-card charge of the original theorem—and the peg is engineered precisely so that nobody has to wait or to discount. The wealth effect this subsection has quantified is therefore not an artefact of our translation but a property of the instrument: whoever builds a par-spendable claim on the bond stock builds the windfall into it.

3.2. Endogenising θ: the issuer's balance sheet

In Section 3.1 the stablecoin entered the model in reduced form: the single parameter θ recorded the share of the bond stock that the payment constraint accepts at par, and Proposition 1 traced the price level to it. Who sets θ, and what determines its size, remained outside the model. We now endogenise the parameter by introducing the institution that carries it: the stablecoin issuer.

The institutional mechanics are straightforward. A user deposits \$100 in outside money with the issuer and receives 100 tokens. The issuer keeps a fraction of the deposit in settlement money and places the remainder in Treasury bills; the funds spent on the bills flow to their sellers and return to circulation. For the questions of this paper, two magnitudes therefore describe the issuer completely: the scale of issuance and the composition of reserves. Let S denote the face value of tokens outstanding, and let $\rho \in [0, 1]$ denote the share of reserves held in settlement money, so that the issuer holds ρS in outside money and debt securities of face value $(1 - \rho)S$. The issuer's balance sheet is then:

Assets	**Liabilities**
Settlement money: ρS	Tokens outstanding: S
Credit claims (T-bills, repo, etc.): $(1 - \rho)S$	

Figure 3: Stylised balance sheet of a stablecoin issuer. Assets consist of settlement money and of debt securities issued by third parties (predominantly Treasury securities and repurchase agreements), the debt securities recorded at face value—the convention fixed when the balance sheet enters the model below; liabilities are the tokens held by users, redeemable at par in outside money.

The modelling strategy is the one Section 2 used for the banking sector, and the parallel is worth making explicit. Banking there is a linear technology $(-q, 1)$, whose conduct is pinned down by competition and no-arbitrage: zero profit determines behaviour, and no further behavioural apparatus is needed. The issuer is modelled in the same spirit. Issuance is a linear transformation—one unit of outside money in, one token out, the proceeds split ρ into settlement money and $1 - \rho$ into bills—and Figure 3 is that technology written as stocks. The two descriptions are equivalent: a technology records a financial arrangement flow by flow, and a balance sheet records the position the flows leave behind. Nothing separates an institution "modelled as a technology" from an institution "with a balance sheet" beyond the notation chosen; the bank of Section 2 could equally be displayed as loan claims held against the inside money created. What does distinguish institutions holding the same assets—a bank, a money market fund, a coin issuer—is the liability side: what the claim each of them issues can do in the constraint system, that is, whether the payment constraint accepts it. That margin is exactly what θ records, and the money-market-fund contrast of Section 3.1 is this observation at work. The linear form also decides where the economics of the institution sits: with behaviour reduced to no-arbitrage pass-through, everything the issuer contributes to the price level is carried by the two parameters the description leaves free—the scale S and the composition ρ—and by nothing else about the institution. Only the payment-relevant side of the institution appears here, by design. The issuer is modelled solely as a transformer of bonds into circulating payment capacity; funding structure, run dynamics, and solvency belong to the financial-stability questions from which the paper abstracts throughout.

The composition parameter ρ has economic content of its own. Interest accrues only on the non-money share of the balance sheet: Circle's S-1 filing revealed that 95–99% of its revenue derives from interest earned on reserve assets (Circle Internet Group, 2025). The margin that makes issuance a business is the same margin that makes reserves debt securities rather than cash—left to itself, the industry chooses ρ close to zero.

Table 3 shows this choice in the data: it reports the reserve composition of the two largest issuers, Tether and Circle, as of September 30, 2025.

The mapping rows at the bottom of the table carry the main information. Settlement media are 0.02 percent of Tether's reserves ($\rho \approx 0.0002$) and 14.1 percent of Circle's ($\rho \approx 0.14$): even the conservatively managed issuer holds roughly six-sevenths of its reserves in debt securities, and the industry leader holds essentially none in settlement money. Two remarks connect the table to the model. First, the row "cash & bank deposits" maps to the model's outside money by way of the consolidation of Section 2: a bank deposit is a claim on the settlement system that has not been deployed into securities, and the compression of the two-tiered architecture treats such claims and central bank liabilities alike. Second, Tether's precious metals and bitcoin fit neither reserve category, and the difference is not one of degree. A Treasury bill is a promise: an obligation of the public sector to deliver settlement money at maturity. A coin backed by bills therefore dissolves, at settlement, into the hierarchy's own money—it is a *derivative* instrument, a *wrapper* around claims to the settlement medium. The term is the industry's own—tokens that repackage an underlying asset are called *wrapped assets*—and we adopt it because the mechanics coincide; it recurs as a term of art below. Gold and bitcoin embody no such promise: nobody is obliged to redeem them into anything, which is what it means for an asset to be nobody's liability. A stablecoin entirely backed by such assets has severed the link to settlement money, and if it circulates, it circulates as a candidate *competitor* to the existing monetary arrangement rather than as a derivative of it. That case is taken up in Section 4.

The table also clears up a point of terminology. A coin backed one hundred percent by Treasury bills is fully reserved in the solvency sense: its assets cover its redemption obligations. But it operates at $\rho \approx 0$ in the sense that matters for the price level: none of its backing is settlement money withdrawn from circulation, and all of it is public debt wrapped into circulating form—a private monetisation of public debt. What regulation and market commentary call the reserve *ratio* is not the relevant object here; the relevant object is reserve *composition*.

The balance sheet also shows what issuance does to the economy's payment capacity, and the mechanics deserve to be traced in full, because the view from inside the acquiring transaction is apt to mislead. Return to the user of the opening example, who gave up one hundred dollars of perfectly good money and received one hundred tokens: from where that

Table 3: Reserve composition of the two largest stablecoin issuers (September 30, 2025). Data from the independently attested Financial Figures and Reserves Report of Tether International, S.A. de C.V. (Tether International, S.A. de C.V., 2025) and the USDC Reserve Report of Circle Internet Group, Inc. (Circle Internet Group, Inc., 2025). The final rows map the reported categories to the model's reserve-composition parameter ρ: settlement media (cash and bank deposits not deployed into debt securities) versus debt securities (assets whose purchase returns funds to the financial system).

	Tether (USDT)		Circle (USDC)	
Asset category	**USD bn**	**%**	**USD bn**	**%**
U.S. Treasury bills	112.4	62.0	21.8	29.5
Overnight reverse repo	18.0	9.9	41.6	56.4
Term reverse repo	3.1	1.7	—	—
Money market funds	6.4	3.5	—	—
Non-U.S. Treasury bills	0.05	0.0	—	—
Corporate bonds	0.01	0.0	—	—
Secured loans	14.6	8.1	—	—
Precious metals (gold)	12.9	7.1	—	—
Bitcoin	9.9	5.4	—	—
Other investments	3.9	2.1	—	—
Cash & bank deposits	0.03	0.02	10.4	14.1
Total reserves	**181.2**	**100**	**73.8**	**100**
Mapping to model parameter ρ				
Settlement media ($\approx \rho$)	0.03	0.02	10.4	14.1
Credit claims ($\approx 1-\rho$)	154.5	85.3	63.4	85.9
Other (precious metals, BTC, etc.)	26.7	14.7	—	—

Tether data from Financial Figures and Reserves Report as of 30 September 2025, attested by BDO Italia S.p.A. under ISAE 3000 (Revised). Circle data from USDC Reserve Report as of 30 September 2025, examined by Deloitte & Touche LLP under AICPA attestation standards. Circle's reserves are held in the Circle Reserve Fund (an SEC-registered 2a-7 government money market fund managed by BlackRock) and in segregated bank accounts. Circle cash includes \$0.8 bn in the Reserve Fund and \$9.6 bn in segregated accounts (both net of timing and settlement differences). Circle's report labels its \$41.6 bn position "U.S. Treasury Repurchase Agreements"; economically it is overnight lending of cash against Treasury collateral, i.e. reverse repo, and it is classified here accordingly. "Credit claims" comprise all assets whose purchase channels funds back into the financial system (T-bills, repo, MMFs, secured loans, corporate bonds). "Settlement media" comprises cash and bank deposits not deployed into debt securities or asset purchases. Tether's precious metals, bitcoin, and other investments are classified separately as they are neither settlement media nor standard debt securities. For Tether, the reported total U.S. Treasury exposure of \$135 bn (cited in the issuer's press release) includes indirect exposure through money market funds and repo collateralised by Treasuries, in addition to the \$112.4 bn in direct T-bill holdings shown here.

user stands, nothing monetary has happened—one liquid instrument has been exchanged for another, and the user's payment capacity is exactly what it was. This introspection is correct, and for the money share of the reserve it is the whole story: the impounded ρS backs tokens that replace it one for one, a neutrality that Proposition 2(i) below states as a theorem. The bond share is different, and the difference is located not at the user but further down the chain of transactions. The money placed against it does not stay with the issuer: net of the issuance margin, it passes through to the sellers of the bills and returns to circulation in their hands, while the tokens it financed circulate in addition. When all positions are recorded—Figure 4 traces them—circulation has lost only the impounded ρS; S tokens that the payment constraint accepts at par circulate on top of what remains; and $(1-\rho)S$ of the public debt stands behind the coin, payable at par where before it was liquid only at the discount q through the credit system. Every party along the chain can truthfully report a swap of equal value for equal value; the aggregate has nevertheless gained payment capacity—the wedge $1-q$ per wrapped bond, exactly the movement the stock channel of Section 3.1 prices. Where the gain lodges—with the bondholder who wraps, or with the issuer as its margin—depends on the path by which the bonds were acquired, a distributional question taken up after the proposition below.

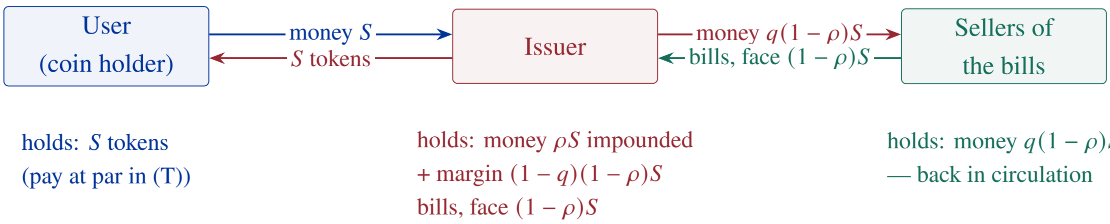

	before	after
outside money in circulation	M	$M-\rho S$
tokens accepted at par in (T)	0	S
public debt spendable at par	0	$(1-\rho)S$

Only the impounded ρS leaves circulation; the rest of the user's money moves on — $q(1-\rho)S$ to the bill sellers, $(1-q)(1-\rho)S$ to the issuer as its margin. The wrapped debt was liquid only at q through the credit system before; it now pays at par.

Net: circulation loses ρS and gains S tokens — a gain of $(1-\rho)S$ in payment capacity, the wedge $1-q$ per wrapped bond.

Figure 4: The issuance chain at $t=0$ and the payment media in circulation before and after issuance, for a coin of scale S and reserve composition ρ. Top: the flows of issuance, arrows colored by the party they emanate from, and each party's position once issuance is complete. Bottom: the aggregate tally. Every party exchanges equal value for equal value, yet payment capacity rises by $(1-\rho)S$—the movement the stock channel of Section 3.1 prices.

The question the balance sheet forces is therefore no longer who sets θ, but what value the issuer's choices imply for it: how much of the bond stock B ends up in form of a circulating coin—θ as a function of the scale S and the composition ρ. Answering it is the task of the remainder of this section.

The first step is a bookkeeping convention with economic content. We record the bond side of the issuer's balance sheet at *face value*: against S tokens outstanding, the issuer holds ρS in settlement money and bonds of face value $(1-\rho)S$. The convention is chosen for what it delivers. At settlement the issuer's resources are the impounded money ρS plus the delivery $(1-\rho)S$ on its bonds—together exactly S, so every token, spent or unspent, redeems at par, and the issuer's books close with nothing left over. And since bonds trade at q at $t = 0$, a balance sheet of this shape is assembled at a margin: face value $(1-\rho)S$ is acquired for $q(1-\rho)S$, leaving $(1-q)(1-\rho)S$ with the issuer—the issuance margin, the model counterpart of the revenue fact documented in the S-1 filing above. One token thus corresponds to one unit of face value, of money or of bonds, which is precisely the correspondence the wrapping device of Section 3.1 uses.[19]

With the convention fixed, the issuer enters the model as an arrangement among households, mediated by its balance sheet. Before trade begins—the same pre-trade timing the coining bound of Section 3.1 imposes—households place money $d^h \in [0, m^h]$ and bonds $w^h \in [0, b^h]$ with the issuer and receive tokens one for one, in quantity $d^h + w^h$; in the aggregate the placements make up the balance sheet, $\sum_h d^h = \rho S$ and $\sum_h w^h = (1-\rho)S$. Feasibility of the balance sheet requires $\rho S \leq M$ and $(1-\rho)S \leq B$: the issuer cannot impound more money, or wrap more bonds, than the economy holds. Tokens enter the payment constraint at par, so household h's payment constraint reads

$$p \cdot \tau^{h,+} + q\beta^h \;\leq\; (m^h - d^h) - q\ell^h + (d^h + w^h) \;=\; m^h - q\ell^h + w^h \;:$$

the deposit cancels before the analysis begins. A money-backed token is a wrapper around an impounded unit of settlement money and replaces it one for one; only the bond-backed share changes what can pay. Because each bond placed gains its holder $1-q$ of payment capacity, every household wants to wrap to capacity, as in Section 3.1; the issuer's wrapping capacity $(1-\rho)S$ must therefore be rationed, and we let it be rationed in proportion to bond holdings: $w^h \leq \theta b^h$ with a common cap θ. Summing the binding caps identifies the cap the balance sheet implies. This is the content of the following result, which is proved in Appendix A by reduction to Proposition 1.

Proposition 2 (Endogenised moneyness)**.** *Fix $q \in (0,1)$ and an issuer of scale $S \geq 0$ and composition $\rho \in [0,1]$ with $\rho S \leq M$ and $(1-\rho)S \leq B$. In the issuer economy just described:*

[19] In Table 3 the distinction between face and market value is immaterial—Treasury bills trade near par—but in the model q lies materially below one and the convention matters. Recording the bond side at its $t = 0$ market value instead would back $(1-\rho)S$ of tokens by face value $(1-\rho)S/q$, breaking the one-for-one correspondence between tokens and wrapped face value and moving the issuer's margin to settlement. Nothing of substance changes; the accounting is merely less transparent.

(i) (Deposited money is a pass-through.) *The household constraint sets depend on the placements* (d^h, w^h) *only through* w^h*: each money-backed token replaces, one for one, the unit of money impounded against it. The money share of issuance leaves payment capacity, prices and the allocation unchanged, household by household.*

(ii) (Induced moneyness.) *The issuer economy has the same monetary equilibria as the coin economy* $\mathcal{E}(\theta)$ *of Proposition 1 at*

$$\theta(S, \rho) \;=\; \frac{(1-\rho)S}{B} \;\in\; [0, 1].$$

In particular, $(1-q)V = M + qB + (1-q)(1-\rho)S$ *in any monetary equilibrium, and under the uniformity hypothesis of Proposition 1(iii) all money prices scale relative to the coinless baseline by*

$$\kappa\big(\theta(S, \rho)\big) \;=\; 1 + \frac{(1-q)(1-\rho)S}{M + qB}.$$

Part (i) upgrades the neutrality valve of the flow-channel discussion below from an assertion to a theorem, and sharpens it: the narrow-bank coin ($\rho = 1$) is neutral not merely in the aggregate but household by household—$\theta(S, 1) = 0$ however large the issuer grows. Part (ii) answers the question this subsection opened with. Who sets θ? Nobody sets it; it is implied. The moneyness of the bond stock is the wrapped share $(1-\rho)S/B$, determined by two magnitudes with identifiable owners: the scale S, set by token demand and licensed by the regulator, and the composition ρ, chosen by the issuer within whatever bounds regulation imposes.

The proposition describes the issuer's bond acquisition in kind: a household places its own bond and receives a token. This, it is worth noting, is the institutional image of Dubey and Geanakoplos (2026) themselves: the money market card of their opening pages—a card that lets its holder "charge on bonds held in a money market" (their §1.1), the arrangement whose adoption question their §6.4 poses—is the in-kind path, a household's own bonds made spendable. The institutional mechanics sketched at the start of this subsection run differently—users deposit money and the issuer buys the bills in the market, which is how the observed issuers of Table 3 in fact operate. The two paths are the same arrangement up to who pockets the margin. Under deposit-and-purchase, the depositor's money passes through the issuer to the bond seller, who parts with one unit of face value per q received; the positions after trade—tokens outstanding, bonds wrapped, money in circulation—are those of the in-kind path, and every aggregate of the proposition is unchanged. What differs is the destination of the wedge: in kind, the wrapping household keeps the gain $1-q$ per bond; under deposit-and-purchase, the issuer keeps it as the issuance margin. The difference is distributional, of exactly the kind the windfall discussion of Section 3.1 prices: invisible to payment capacity, trade volume and the absorption identity, visible to relative prices and the allocation only through the skew of who gains. The equivalence is thus also what connects the framework to the institutions: Dubey and Geanakoplos' charge card and the deposit-and-purchase issuer of Table 3 are one arrangement in every respect that matters for the price level.

The formula for κ shows what each of the two magnitudes does to the price level. Scale and composition enter through the single product $(1-\rho)S$, so the price factor is linear in S and linearly throttled by ρ: each additional token adds $(1-\rho)(1-q)$ to payment capacity—nothing through its money-backed share, the full wedge through its bond-backed share. In the running example ($q = 0.6$, $B = M/2$), an issuer wrapping the entire bond stock at $\rho = 0$ reproduces the 15.4-percent price rise of Section 3.1; the same scale at Circle's measured $\rho \approx 0.14$ yields 13.2 percent; at $\rho = 1$, nothing. Read together with the incentive observation above, the formula also states the political economy of the sector in one line: the margin $(1-q)(1-\rho)S$ the issuer earns and the addition $(1-q)(1-\rho)S$ the price level must absorb are the same number. The business model and the inflation channel are one object read from two sides—which is why the industry, left to itself, gravitates to the corner where both are largest: ρ near zero, and S as large as the demand for tokens permits. What limits that demand is the breadth of the coin's acceptance, the final determinant of the taxonomy, taken up at the end of the section.

3.3. The flow channel: when new debt arrives spendable

The two preceding subsections held the stock of public debt fixed. The coin changed what a given stock B could do—first by assumption, through the parameter θ, then through the issuer's balance sheet that carries it. This subsection asks what happens when the stock itself responds: stablecoin growth creates demand for Treasury bills as reserve assets, and the fiscal authority may meet that demand with new issuance.

It is best to state at the outset what this channel is not. That more public debt raises the price level is a property of the baseline economy of Section 2, established there before any stablecoin exists: bonds are outside wealth, and adding them to endowments raises nominal prices—Dubey and Geanakoplos' own comparative static (Dubey and Geanakoplos, 2026), rehearsed numerically at the close of Section 2. Nor is there anything specifically monetary about the purchase itself: whoever buys the new bills—a money market fund, a pension fund, a stablecoin issuer—absorbs the issuance in the same way. The flow channel is therefore not a second mechanism alongside the stock channel; it is a movement *along* a relation the baseline already contains. What is specific to the stablecoin is the economy into which the marginal bond arrives. Where a coin sector stands ready to provide the financial engineering to create concurrent payment capacity from government bonds, a newly issued bond does not wait for the credit system to monetise it at the discount q: a fraction θ arrives spendable at par.[20] The same fiscal act moves prices further—and by a factor the model makes exact.

Whether the issuer's demand for bills is met by new issuance is a choice of the fiscal authority, not of the issuer. The Treasury manages the maturity structure of the public debt in response to demand, and the bill margin is the elastic one: short-term issuance is the standard instrument for absorbing shifts in the demand for safe assets. Regulation has meanwhile hard-wired the connection on the demand side: by restricting the interest-bearing part of permissible reserves to precisely this segment—bills and Treasury-collateralised repo—the GENIUS Act ensures that every dollar of coin growth the issuer deploys for yield arrives as demand in the bill market. What the Treasury does with that demand, however —accommodate it with new bills or let it bid for the existing stock—is a fiscal decision the model cannot derive. It must be assumed.[21]

Assumption 1 (Fiscal response)**.** *Of the face value of the bonds the issuer acquires, a fraction $\eta \in [0, 1]$ is met by net new issuance; the remaining fraction $1 - \eta$ is purchased from existing private holders. New issuance enters the economy as an addition to endowed bond wealth: the bond stock rises from B to $\tilde{B} = B + \Delta B$.*

The formulation divides the work deliberately between what is assumed and what is proved. The assumption fixes only a quantity—how much new debt is created—because this is fiscal behaviour, which the model cannot predict. What the new debt then does is proved, not assumed: bonds are outside wealth in this economy, so a larger bond stock raises payment capacity and the price level—the content of the corollary below. The division is worth a comment because it is not available in every framework. In a model without endowed bonds, a parameter in η's position would have to carry a second assumption on top of the quantity: that government deficits add to private net wealth at all (the non-Ricardian step). Here that step is a result of the Dubey–Geanakoplos framework, not an assumption of ours. And reading the endowment change as a deficit is licensed by Section 2: endowments are compressed fiscal history, so ΔB is simply new expenditure not taxed back.

[20]"Arrives spendable" is to be read in the comparative-static sense in which the model speaks. New issuance is a change of endowments between two economies, and in the economy with the enlarged stock the endowed bonds—new and old alike—are spendable in the proportion θ. Within any single trading date, the coining bound of Section 3.1 continues to apply: coins are minted against endowments before trade begins, so a bond *purchased* at price q during the date can never be spent at par—the arbitrage that bound exists to exclude.

[21]The demand side of the nexus is public record: on the GENIUS Act's passage in the Senate in June 2025, the Treasury Secretary welcomed precisely the prospect that a thriving stablecoin sector "will drive demand from the private sector for US Treasuries" (Bessent, 2025). The statement documents that the issuance response is a live policy scenario; whether the demand is met with new bills—the parameter η below—it does not settle.

Corollary 2 (Fiscal amplification)**.** *Fix $q \in (0,1)$ and $\theta \in [0,1]$, and consider the economy $\mathcal{E}(\theta)$ of Proposition 1 with bond stock B.*

(i) In any monetary equilibrium, nominal payment capacity satisfies $(1-q)V = M + qB + \theta(1-q)B$: a marginal unit of bond face value contributes $q + \theta(1-q)$, against q in the absence of a coin sector. The amplification factor is $1 + \theta(1-q)/q$, and the cross effect is $\partial^2\big[(1-q)V\big]/\partial B\, \partial\theta = 1 - q > 0$.

(ii) If bond endowments before and after issuance satisfy the uniformity hypothesis of Proposition 1(iii), the enlarged stock $\tilde{B}$ supports a monetary equilibrium with the same allocation and all money prices scaled by $\big(M+q\tilde{B}+\theta(1-q)\tilde{B}\big)/\big(M+qB+\theta(1-q)B\big)$.

The proof is a re-application of Proposition 1 to the enlarged endowments and is given in Appendix A. The numbers of the running example make the amplification concrete: at $q = 0.6$, a marginal bond moves payment capacity by 1.0 per unit of face value when $\theta = 1$, against 0.6 when $\theta = 0$—the same deficit is $5/3$ as inflationary—and by the factor $4/3$ at $\theta = \frac{1}{2}$. This is the precise content of the interaction anticipated at the end of Section 3.1: deficits become more inflationary per bond issued as θ rises.

One reading note connects the corollary to the endogenisation of Section 3.2. There $\theta = (1-\rho)S/B$, so holding θ fixed while B grows implicitly assumes that the wrapped face value $(1-\rho)S$ grows in proportion to the stock—the coin sector expanding pari passu with the debt. The composed comparative static at the end of this subsection drops this assumption and lets the scale S and the stock B move separately.

What happens to bank credit deserves its own paragraph, because the natural guess—that the new payment capacity crowds out borrowing—has the sign wrong. The settlement absorption condition of Corollary 1 pins aggregate credit: $(1-q)\bar{D} = M + \tilde{B}$, so equilibrium credit rises with the stock, by $\Delta B/(1-q)$—two and a half units per unit of new debt at $q = 0.6$. Far from being crowded out, inside money *expands* to absorb the larger settlement obligation at the prevailing policy rate, exactly as baseline credit expands when M rises at fixed q. The amplification of Corollary 2 and the credit expansion are not competing channels but the same equilibrium adjustment seen from the payment side and from the settlement side.

Two limiting cases bound the channel—two shut-off valves. At $\rho = 1$ the issuer holds only settlement money: every token wraps an impounded unit of outside money, payment capacity is unchanged, no demand reaches the bill market, and, by Proposition 2, $\theta(S, 1) = 0$. The narrow-bank coin is neutral. At $\eta = 0$ the fiscal authority does not respond: coin demand bids for the existing stock, B is unchanged, and the flow channel is shut, while the stock channel of Proposition 1 operates unimpaired. Each valve is a policy lever—reserve composition for the first, debt management for the second—and both return in Section 6.

These statements can be checked computationally. Table 4 enlarges the bond stock of the running example from $B = 1.1$ to 1.65 and 2.2 at fixed $\theta \in \{0, \frac{1}{2}, 1\}$, scaling bond endowments proportionally so that the uniformity hypothesis of Corollary 2(ii) is preserved. The predictions are borne out exactly: the marginal contribution of a unit of bond face value to payment capacity is 0.6, 0.8 and 1.0 at $\theta = 0, \frac{1}{2}, 1$—the corollary's $q + \theta(1 - q)$ to solver precision—prices scale by the predicted factor, the allocation does not move, and all three equilibrium identities hold at the tolerance of the computation. Two readings of the same numbers deserve emphasis. The contribution column is the amplification made visible: a unit of new debt adds 0.6 to payment capacity in the unwrapped economy and 1.0 in the fully wrapped one—the marginal bond is 5/3 as potent, and doubling the stock accordingly raises the price level by 23.1 percent at $\theta = 0$ against 33.3 percent at $\theta = 1$. And within each $\tilde{B}$ column, aggregate credit $\bar{D} = (M + \tilde{B})/(1 - q)$ is identical across θ: the credit expansion of the absorption condition responds to the stock alone, while θ decides how much trade the same settlement obligations finance.

Table 4: Fiscal amplification in the running example. Each block fixes the spendable share θ and enlarges the bond stock from $B = 1.1$ to $\tilde{B} \in \{1.65, 2.2\}$, scaling bond endowments proportionally so that the uniform portfolio composition of the running example ($b^h = m^h/2$) is preserved and Corollary 2(ii) applies. Equilibria are computed via the equivalence of Proposition 1(i) (endowments $m^h + \theta\tilde{b}^h$, $(1 - \theta)\tilde{b}^h$; dgme package v0.4.0). Parameters: $q = 0.6$, $M = 2.2$. scale: observed price factor $p_l(\tilde{B})/p_l(B)$ against the predicted $\big(M + q\tilde{B} + \theta(1 - q)\tilde{B}\big)/\big(M + qB + \theta(1 - q)B\big)$; contribution: observed marginal contribution of a unit of bond face value to payment capacity, $(1 - q)\big(V(\tilde{B}) - V(B)\big)/(\tilde{B} - B)$, against the predicted $q + \theta(1 - q)$ of Corollary 2(i); V: nominal trade volume; $\bar{D}$: aggregate credit; max $|\Delta x|$: largest change in any household's consumption relative to the $\tilde{B} = B$ row of the same block; residual: largest absolute violation of the three equilibrium identities $(1 - q)V = M + q\tilde{B} + \theta(1 - q)\tilde{B}$, $\bar{D} = V + (1 - \theta)\tilde{B}$ and $(1 - q)\bar{D} = M + \tilde{B}$.

θ	$\tilde{B}$	scale (obs / pred)		contribution (obs / pred)		V	$\bar{D}$	max $\lvert\Delta x\rvert$	residual
$\theta = 0.00$ *(amplification factor* $1 + \theta(1 - q)/q = 1.000$*)*									
0.00	1.10	1.0000	1.0000	—	0.60	7.15	8.25	0.000	$4.5e - 12$
0.00	1.65	1.1154	1.1154	0.6000	0.60	7.97	9.62	0.000	$5.2e - 12$
0.00	2.20	1.2308	1.2308	0.6000	0.60	8.80	11.00	0.000	$5.9e - 12$
$\theta = 0.50$ *(amplification factor* $1 + \theta(1 - q)/q = 1.333$*)*									
0.50	1.10	1.0000	1.0000	—	0.80	7.70	8.25	0.000	$5.1e - 12$
0.50	1.65	1.1429	1.1429	0.8000	0.80	8.80	9.62	0.000	$5.9e - 12$
0.50	2.20	1.2857	1.2857	0.8000	0.80	9.90	11.00	0.000	$6.6e - 12$
$\theta = 1.00$ *(amplification factor* $1 + \theta(1 - q)/q = 1.667$*)*									
1.00	1.10	1.0000	1.0000	—	1.00	8.25	8.25	0.000	$5.3e - 12$
1.00	1.65	1.1667	1.1667	1.0000	1.00	9.62	9.62	0.000	$6.4e - 12$
1.00	2.20	1.3333	1.3333	1.0000	1.00	11.00	11.00	0.000	$7.3e - 12$

The endogenisation and the fiscal response now compose into a single comparative static, worth recording because it prices the marginal coin dollar. Let the sector grow by ΔS at composition ρ, and let the Treasury respond as in Assumption 1. The wrapped face value rises by $(1-\rho)\Delta S$, the bond stock by $\eta(1-\rho)\Delta S$, and payment capacity by

$$\Delta\big[(1-q)V\big] \;=\; (1-\rho)\,\Delta S\,\big[(1-q)+q\eta\big]\;:$$

the bond leg of each new token adds the wedge $1-q$ by wrapping, and, to the extent the Treasury accommodates, a further q through the new bond's own addition to outside wealth. At $\eta = 0$ only the stock channel operates: in the running example, the marginal coin dollar at $\rho = 0$ adds 0.4 to payment capacity. At $\eta = 1$ it adds the full 1.0: a bill issued directly into the issuer's reserve arrives already wrapped—deferred money made spendable on the day it is issued. Between the valves, the expression locates every actor of the taxonomy in one line: ρ belongs to the issuer and its regulator, η to the fiscal authority, q to the central bank, and S to the market.

3.4. Adoption: breadth and speed

The channels of this section share a final determinant, so far held at its maximum. Throughout, the coin has been allowed to pay for *all* goods: the acceptance leg of the opening scope remark was assumed not only present but complete. Adoption breadth relaxes this. Suppose sellers of only a subset A of the goods accept the coin, while the remaining goods trade against money alone.[22]

The consequence can be stated without new apparatus. A token pays only inside A, while an unwrapped bond monetises at q towards *all* goods; a household therefore wraps only what its purchases in the accepting set can use, and tokens the accepting set cannot absorb are not demanded in the first place. The effective wrapped stock is capped by the trade that can absorb it:

$$\theta_{\text{eff}}\,B \;=\; \min\big\{\,(1-\rho)S,\ \text{nominal coin-payable trade in } A\,\big\},$$

and the stock channel operates at θ_{eff}, not at the engineered $\theta(S,\rho)$ of Proposition 2.

We state this as a comparative-static sketch rather than a theorem, and the status of the two objects in the display should be explicit, because neither is a parameter anyone can set. Take the accepting set A first. Which sellers accept the coin is itself an economic outcome—the acceptance decision of the two-sided market in the scope remark that opened the section—and deriving A inside the model would require a demand system in which

[22] The predecessor of this device in an earlier version of this paper was a three-way partition of transactions into money-only, token-only and mixed categories. The single accepting set suffices for the point made here; the token-only category returns below as the reading of present-day circulation.

sellers choose what to accept, and an equilibrium concept under which these choices are mutually consistent. We do not build this apparatus, and nothing in the paper needs it: A enters as a description of the acceptance infrastructure in place, the model's given. But even with A given, θ_{eff} is not free: the bound's second argument, nominal coin-payable trade in A, is evaluated at the equilibrium prices that the coin's own circulation helps determine. The display is therefore a consistency condition an equilibrium must satisfy, not a formula over primitives. Making it one would split the payment constraint into an accepting and a non-accepting part—exactly as the two-unit system in the proof of Proposition 3 splits it between currencies—and would add nothing to the economics of the bound. What the bound says is that the two legs of the two-sided market are the two arguments of a minimum: financial engineering supplies moneyness, acceptance absorbs it, and the channel runs at the smaller of the two.

The bound also gives the empirical honesty of the opening scope remark its model form. Today the accepting set of ordinary goods trade is close to empty: coins circulate within crypto-asset markets and payment-industry back-ends, not at the checkout (BIS, 2025; Adrian et al., 2025). In the bound's terms, the engineering leg exists at scale—hundreds of billions of face value stand wrapped—while the acceptance leg holds θ_{eff} near zero. The channel is in place; its cap is at zero. This is also the precise sense in which the acceptance leg is a policy margin: every step that widens A—merchant integration, card-network settlement of coin balances, the legal and technical ease of taking tokens in payment—raises the cap directly, moving θ_{eff} towards the engineered $\theta(S, \rho)$. The amplify posture flagged in the footnote of Section 3.1 operates on exactly this margin.

Breadth is a comparative static; its change over time is not. The model compares economies and is silent on the path between them—adoption speed lies outside the two-date frame—so here we only report how Dubey and Geanakoplos read it. A money innovation, they write, "creates a one shot rise in the price level (perhaps spread over several years as the new money grows), but will stop on its own when the new money stops growing" (Dubey and Geanakoplos, 2026, §1.9). The stock channel's level effect would accordingly appear in the data not as a jump but as a drawn-out drift while θ_{eff} climbs towards $\theta(S, \rho)$—a drift easily attributed to other causes, which is the scenario of misread inflation their lecture warns of, and their argument for measuring the new monies directly rather than inferring them.

The taxonomy announced at the opening of the section is now complete, and each determinant has an owner. The moneyness of the bond stock is engineered by the issuer's scale and composition (S, ρ) and licensed by whoever regulates them; the stock B and the issuance response η belong to the fiscal authority; the wedge $1 - q$ to the central bank; and the breadth that decides how much of the engineered moneyness becomes effective belongs to whoever builds—or declines to build—the acceptance leg. What remains is a response and an escalation: the escalation in which the coin outgrows the taxonomy altogether (Section 4), and the question of what monetary policy can do within it (Section 5).

4. Escalation: Stablecoins as competing outside money

The preceding analysis treats stablecoins as derivative payment instruments—wrapped claims that resolve into sovereign outside money at settlement.[23] Their value is pinned down by the redemption commitment; as long as that commitment is credible, stablecoins inherit the price stability of the sovereign currency. The inflationary channels we discussed in Section 3 raise the price level within a nominally anchored system. The escalation we now discuss attacks the anchor itself. What is at risk is not a larger movement of the price level but the nominal determinacy that made the comparative statics of Section 3 well-defined.

The two sections should accordingly be weighed differently. The channels of Section 3 are bounded comparative statics, and a reader can accept every one of them and still discount their near-term relevance on adoption grounds: while ordinary retail acceptance remains thin, the effective moneyness of the bond stock stays near zero (Section 3.4). That discount does not dispose of the present section. The circuit described below does not require the coin to be widely used in retail payments first: it requires receipts that are held rather than redeemed, and obligations contracted and discharged inside the token—a pattern that crypto-asset markets already exhibit and that token-denominated credit would generalise. And what erodes as the circuit consolidates is not a price but an institution: the settlement hierarchy that gives the price level its anchor, an object that appears on no balance sheet and whose loss no market prices in advance. The developments this section describes are further from today's data than those of Section 3.

This scenario is not purely hypothetical. The GENIUS Act provides the legal framework for T-bill-backed stablecoins at scale; payment networks are integrating

[23] The dividing line between money whose value is anchored in a claim on the issuer and money whose value rests on resale expectations alone is drawn with particular clarity by Hellwig (1985), who shows that the case for currency competition stands or falls with it. Our derivative/competitor distinction is that line, transposed into the settlement hierarchy. We also adopt his use of the term: *inside money* comprises every private payment instrument that gives its bearer a claim on the issuer and thereby remains subordinate to the settlement hierarchy—bank credit in the exposition of Section 2, redeemable coins here. In Gurley and Shaw (1960), where the terminology originates, the line is drawn by the backing instead: inside money is matched by private debt and nets out within the private sector. For bank money the two criteria coincide. The wrapper coin is the instrument on which they come apart—privately issued and redeemable, hence inside in Hellwig's sense, yet backed by claims on the public sector—which is why it is derivative in the present section's sense and nevertheless non-neutral in Section 3.

stablecoin settlement; and the conversion plumbing that lets a recipient of tokens end up with ordinary money—the leg without which a treasury could not spend token-denominated receipts—has already been piloted at retail scale in Singapore.[24] The institutional scaffolding for a parallel monetary circuit—one in which stablecoins circulate, settle transactions, and support their own credit creation without returning to outside money for final settlement—is taking shape.

We identify three stages of escalation, described here in institutional terms. The analysis that follows then locates each stage in the model's constraint system—an instrument's position in the settlement hierarchy is defined by which constraint accepts it, payment (T) or settlement (R)—and derives what can be derived: the positions, their boundary, and their end-states. The transitions between the stages are institutional developments on which the model is silent. Section 4.3 closes with an explicit inventory of which elements of this section are results and which are institutional judgments.

Stage 1: The current regime under the GENIUS Act. Stablecoin issuers collect outside money, invest reserves in Treasury bills, and issue tokens redeemable at par. Stablecoins are payment instruments; outside money is the settlement asset. In the model's terms: the coin enters the payment constraint (T) alone, every obligation still discharges in outside money at (R), and the price-level effects are the bounded channels of Section 3.

Stage 2: the self-sustaining circuit. The GENIUS Act (U.S. Congress, 2025) in effect designates short-term Treasury securities as the canonical reserve asset for stablecoin issuers: its reserve menu also admits cash, insured deposits and central bank balances, but the interest-bearing part is confined to bills and Treasury-collateralised instruments. This creates a powerful alignment of incentives. The Treasury obtains a captive demand base for short-term debt: every dollar of stablecoin issuance generates demand for T-bills. The issuer obtains a legally sanctioned reserve with a yield spread—the seigniorage that drives its business model. Crucially, the Act does not regulate the Treasury's issuance behaviour; it regulates what issuers may *hold*. Nothing prevents the Treasury from issuing bills directly against stablecoin demand, or from welcoming the expansion of a buyer base that lowers its funding costs.

[24] In the pilots of Project Orchid, run by the Monetary Authority of Singapore with the government's digital-services agency and industry partners, government vouchers were disbursed as programmable tokens: recipients spent the tokens at participating merchants, and the issuing bank unwrapped the tokens and credited the merchants' bank accounts in ordinary money at the end of each day—merchants never handled a token at all. A parallel industry trial wrapped the regulated XSGD stablecoin the same way. See Monetary authority of Singapore (2022). The purpose of the pilots was the efficiency of fiscal transfers, and their operators were the central bank and the government's own technology agency; a change of monetary order was on nobody's mind. That is exactly what the example teaches: the conversion plumbing a parallel circuit would need is being built, competently and for good reasons, as ordinary payment innovation.

For the Treasury to benefit from this arrangement, it must be able to convert the stablecoin proceeds it receives into usable funds. This requires either that stablecoins are generally accepted for government payments or—more immediately—that intermediaries exchange stablecoins for outside money on the Treasury's behalf. The latter infrastructure already exists: the Singapore pilots described above converted token-denominated government disbursements into bank money at the merchant level, and nothing in the GENIUS Act's licensing framework prevents comparable arrangements from emerging in the United States. The operational obstacle to a direct Treasury–stablecoin channel is therefore not technological but a matter of policy.

As the stablecoin sector grows, a payment infrastructure develops around the tokens—the acceptance leg that Section 3 took as given is here being built. Merchants accept stablecoins; sellers hold them rather than redeeming into dollars; transactions settle stablecoin-to-stablecoin on distributed ledgers. The dollar redemption right exists legally but is exercised less and less frequently. The circuit becomes operationally self-sustaining: tokens circulate, credit is extended against token-denominated collateral, and obligations contracted *within* the circuit are discharged in tokens. The decisive boundary, however, still holds: bank debt, taxes, and final settlement—the economy-wide settlement constraint (R)—continue to resolve in outside money. In the model's terms, Stage 2 is a large θ with a dormant redemption right, not yet a new settlement asset.

At this stage, neither the Treasury nor the issuer has a commitment device to stop the expansion. The Treasury benefits from cheap funding; curtailing issuance means giving up a demand base. The issuer benefits from the spread; curtailing issuance means giving up revenue. And the payment infrastructure, once built, cannot easily be dismantled. The monetary expansion becomes *hard to reverse ex post*: even if the authorities recognise the inflationary consequences, the institutional configuration that produces them has become self-reinforcing, and no party inside it retains an incentive to unwind it.[25]

Stage 3: competing outside money. The final stage is reached when the stablecoin circuit no longer depends on outside money for settlement. Tokens function as their own settlement medium; the redemption link to dollars becomes a legal fiction maintained for regulatory compliance rather than an operational constraint on the monetary system. In the model's terms, the coin has crossed the decisive boundary: it has migrated from the payment constraint into the settlement constraint, and (R) itself now accepts tokens. In the terminology of Borio et al. (2026), the stablecoin issuer has acquired the "exorbitant privilege" of payment elasticity: the capacity to purchase assets by issuing its own

[25]We are grateful to Paul Pichler for this observation and for identifying the incentive alignment under the GENIUS Act that makes Stage 2 a natural consequence of the current regulatory trajectory.

liabilities, rather than first collecting outside money through coin sales.[26] At this point, the economy has two outside monies, and the formal analysis of the present section applies.

Note that none of this requires design. The scenario needs no constituency working toward de-anchoring: each step is individually rational and mostly benign, while the anchor those steps jointly erode appears on no balance sheet and in no objective. The risk is less sophisticated intent than naive financial engineering—instruments built as if the monetary order were a free and indestructible input. Figure 5 summarises the three stages.

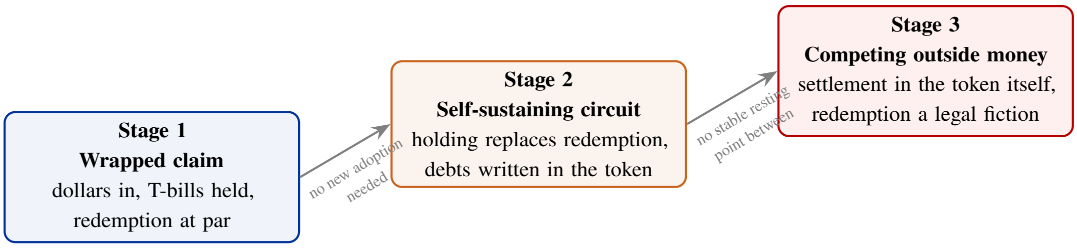

Figure 5: The escalation ladder. Stage 1 is the current regime: the coin is a wrapped claim that resolves into outside money at settlement, living in the payment constraint (T). Stage 2 needs no new adoption, only holders who keep the coin and debts written in the token; redemption falls dormant while economy-wide settlement still resolves in outside money. At Stage 3 the token has entered the settlement constraint (R) and circulates as a second outside money. The second step admits no stable resting point in between: a coin that survives the loss of its redemption anchor must cross into settlement (Section 4.1). No stage requires intent.

The dividing line the stages cross is not foreign to Dubey and Geanakoplos, though in the lecture it appears only in passing and in a different guise: as a taxonomy of what *backs* the new monies. Coins wrapped around nominal claims—Treasuries, credit-card receivables—are, they observe, "themselves dependent on the old money": every such instrument resolves into dollars at settlement, so it can raise the price level but cannot displace the money it settles into (Dubey and Geanakoplos, 2026, §1.8). Coins backed by real assets—gold, equity indices—have severed that dependence, and there their assessment changes sharply: such monies "threaten the viability of old fashioned money" (their §1.8); their §6.6 shows—in a proposition titled "Commodity Monies Can Destroy the Dollar"—that with enough goods trading directly against goods, cash loses its value and monetary equilibrium fails to exist; and the closing scenario of their §1.9 attaches the word hyperinflation to precisely this case. The mechanism carrying their result is specific to real backing: a commodity money opens goods-for-goods markets through which trade can

[26] Borio et al. (2026) see stablecoin issuers sharing banks' privilege of paying with their own liabilities once the coin is widely accepted as a means of payment. The constraint system suggests a two-step reading. Wide acceptance alone—Stage 2—gives the issuer at most what banks have: elasticity disciplined by redemption, every token issued returning to the issuer as a claim to be settled in outside money. What makes the privilege exorbitant in the original sense—financing without that discipline—is the disappearance of the settlement obligation itself: acceptance of the token at (R). In the vocabulary of the stages, banks' kind of privilege can arrive with Stage 2; the exorbitant one is a property of Stage 3, not of scale.

bypass money altogether, eating away the gains from trade that sustain money's value. The settlement-hierarchy frame of this paper simplifies the observation and extends its reach. What matters in the stages above is not that the backing is physical but that the instrument need not return to outside money at settlement: the boundary runs through the constraint system, between what only (T) accepts and what (R) accepts. A Treasury-backed coin sits on the safe side of the backing taxonomy—its reserves are nominal claims, the case Dubey and Geanakoplos class as dependent—and Stage 3 shows that it can nevertheless cross the same boundary institutionally, once redemption falls dormant and settlement practice absorbs the token. The end-states differ in form: in their commodity case monetary equilibrium is destroyed, while in the two-money economy analysed below it survives as a continuum without an anchor. But the underlying dichotomy is one and the same, and the constraint system states it in a single line: what an instrument does to the monetary order is decided at the settlement constraint, not by what stands behind the token.

4.1. Why severing redemption forces a settlement role

Recall why sellers accept the coin at par in the economies of Section 3: all sales proceeds, money and coin alike, become available only at settlement, where the coin delivers exactly one unit of outside money. Acceptance was derived from the delivery. This observation has a converse, and the converse is the escalation's hinge. Augment the two-date economy by an additional instrument that households may tender in payment at $t = 0$, and suppose the instrument neither delivers outside money at settlement nor is accepted in discharge of settlement obligations. Then in any monetary equilibrium its price is zero. We call this the *anchoring dichotomy*, and we state it without the dress of a numbered result because the argument is a single observation. Sales proceeds accrue at settlement, and the economy ends there: at $t = 1$ an object has value only through what it can do in the settlement constraint—discharge an obligation directly, or deliver the outside money that does. An instrument that can do neither is worthless at $t = 1$; a seller who surrendered goods for it at $t = 0$ would strictly reduce her attainable consumption, contradicting optimality at any positive price. This is the two-date face of the Hahn problem (Dubey and Geanakoplos, 2026, §3.5): the value of an intrinsically worthless instrument must be anchored at the horizon by an obligation that demands it.

The consequence for the escalation is that there is no stable resting point between the stages. A coin whose redemption link has been severed either becomes worthless—the fate the dichotomy assigns—or it retains value because settlement obligations have come to be dischargeable in it: because it has crossed into (R). Severing redemption thus does not create a floating private money: the coin either loses its value or crosses into settlement. The narrative's final stage is not one scenario among several but the only destination consistent with the coin surviving the loss of its anchor. We therefore analyse the destination.

4.2. Two outside monies: Existence and indeterminacy

The Stage-3 economy is the baseline of Section 2 with a second outside money: households hold endowments (m^h, b^h, m_S^h), where m_S^h denotes coin holdings with aggregate $M_S > 0$—the legacy stock of the coin era, circulating without redemption, just as (m^h, b^h) are the sediment of past fiscal history. Let $\varepsilon > 0$ denote the exchange rate (coin per unit of outside money), so that coin prices satisfy $p^S = \varepsilon p$ by no-arbitrage. Two assumptions define the stage. First, households can convert freely between the two units at rate ε within the trading date; with free conversion it is immaterial in which unit bank credit is denominated, and we let the bank lend in sovereign units at the price q. Second—the defining feature—settlement obligations can be discharged in either unit at the same rate: the coin is accepted in (R). Both assumptions are strong, and deliberately so; they constitute the frictionless benchmark, and we return below to what they mean.

A *monetary equilibrium with two outside monies* consists of prices $(\bar{p}, q, \bar{\varepsilon})$ with coin prices $\bar{p}^S = \bar{\varepsilon}\bar{p}$, and household plans that are optimal in the two-unit budget set—payments constrained by holdings of each unit after conversion, settlement obligations dischargeable in either unit at $\bar{\varepsilon}$—together with goods-market clearing and accommodation of aggregate credit at the policy rate q. The formal constraint system is stated with the proof in Appendix A.

Proposition 3 (Existence and indeterminacy with two outside monies)**.** *Suppose $\gamma(e) > r$, and let coin holdings be proportional to liquid wealth: $m_S^h/M_S = (m^h + qb^h)/(M + qB)$ for all h. Then:*

(i) for every $\bar{\varepsilon} > 0$ there exists a monetary equilibrium with two outside monies;

(ii) the exchange rate $\bar{\varepsilon}$ is not determined by the equilibrium conditions;

(iii) the equilibria of part (i) can be chosen so that the allocation is the same for every $\bar{\varepsilon}$, while all sovereign-denominated prices scale with the effective liquid wealth $M + M_S/\bar{\varepsilon} + qB$ and coin-denominated prices equal $\bar{\varepsilon}$ times sovereign prices.

The proof is by aggregation. Fix $\bar{\varepsilon}$ and define sovereign-equivalent holdings $\tilde{m}^h \equiv m^h + m_S^h/\bar{\varepsilon}$. Free conversion and free discharge collapse the two payment and settlement constraints into the single constraint system of Section 2 at the aggregate $\tilde{M} = M + M_S/\bar{\varepsilon}$; the collapse lemma and the existence theorem deliver an equilibrium of the aggregated economy, which decomposes into a two-money equilibrium at the chosen $\bar{\varepsilon}$. Since every equilibrium condition involves the two money stocks only through the aggregate $\tilde{M}$, no condition restricts $\bar{\varepsilon}$: the equilibrium set is a continuum indexed by the exchange rate. Proportionality makes each household's share of effective liquid wealth invariant to $\bar{\varepsilon}$, which delivers the constant allocation and the uniform price scaling, by the argument of Corollary 2(ii). Details are in Appendix A.

The construction is consistent with the anchoring dichotomy of Section 4.1, and it is worth pausing to see how. The legacy stock M_S circulates without redemption—exactly the kind of instrument the dichotomy declares worthless. It escapes that fate here only because Stage 3's defining feature grants what the dichotomy requires: settlement obligations accept the coin, so (R) itself anchors it. The two results are therefore the two horns of one statement, not in tension. Strip the coin of its settlement acceptance and its value collapses to zero; grant it, and positive value becomes possible, but nothing pins its level—the continuum of the proposition.

The result is the finite-horizon, settlement-based expression of the classic indeterminacy of Kareken and Wallace (1981): when two intrinsically worthless monies are both acceptable and freely convertible, no equilibrium condition pins their relative price. The mechanism here is instructive in one respect. Kareken and Wallace obtain the indeterminacy in an overlapping-generations economy, where money is valued through an infinite chain of expectations; one might suspect the pathology of living in that chain. Here the horizon is finite and money is valued through settlement—and the indeterminacy appears all the same, because once both units are accepted in (R) and freely convertible, nothing in the settlement structure distinguishes them. The loss of determinacy is not an artifact of the infinite future; it is the direct consequence of surrendering the settlement margin.

The discussion of currency competition by Hellwig (1985) contains the sharpest illustration of what this indeterminacy means—and of what ordinarily prevents it. Consider, with Hellwig, two "currencies": DM 100 notes and DM 50 notes. That one trades for two of the other looks like arithmetic; it is in fact a market price, sustained by the expectation that it will hold in future transactions—and any other fixed ratio would equally support a rational-expectations equilibrium. What ties the market to the arithmetic is, in part, the issuer: the Bundesbank stands ready to exchange the denominations at par. A pegged stablecoin is exactly such an intervention-backed parity, with the redemption commitment as the standing intervention—and Stage 3 is the removal of the intervention. The indeterminacy of Proposition 3 is thus not an exotic possibility engineered by the model; it is the generic condition of coexisting outside monies, from which anchored systems are protected only by the commitments that Stage 3 dissolves.

Part (iii) is the central statement of this section. The indeterminacy is not confined to the relative price of the two monies: sovereign prices scale with $M + M_S/\bar{\varepsilon} + qB$, so the *sovereign price level itself* varies along the equilibrium continuum. Stage 3 does not add a volatile asset price to an otherwise anchored system; it removes the fact that made the price level a well-defined object. This is what "attacking the anchor" means, stated as a theorem rather than a metaphor.

4.3. Economic interpretation

The indeterminacy result merits careful interpretation. Under the frictionless benchmark—free conversion, discharge in either unit, a single credit price—the two-money economy is, by Step 1 of the proof, a re-denomination of a single-money economy: the exchange rate scales nominal magnitudes without touching real ones. One might object that this makes the result a numeraire artifact. But the economics is precisely the point: absent an institutional anchor, the exchange rate between two outside monies has nothing fundamental to pin it down. In the frictionless benchmark this surfaces as a continuum of equilibria; with frictions—conversion costs, settlement delays, counterparty risk—the indeterminacy does not disappear but reappears as a range of sustainable rates, or as multiple equilibria selected by expectations and coordination. The frictionless case is the cleanest analytical statement of a genuine economic problem: competing outside monies lack a natural parity. And the observation cuts both ways: it is exactly conversion frictions and mandatory settlement in sovereign money—the hierarchy—that protect the anchored system from the continuum. Stage 3 is their dissolution.

The proposition, moreover, understates the problem, because it holds the coin stock M_S fixed: it is a statement about valuation, not about supply. Hellwig (1985), in his critique of the Hayekian proposals for unregulated currency competition (Hayek, 1976), supplies the other half. Among freely convertible outside monies there can be no systematic differences in inflation rates—any expectation that one money will be relatively worthless tomorrow makes it relatively worthless today—so the disciplining mechanism that competition is supposed to provide has nothing to work on. Competition among suppliers fares no better: price-taking or oligopolistic supply of an outside money drives its value to zero, and an unregulated monopolist over-inflates for seigniorage. Deepest of all, without binding precommitments about future supply, no equilibrium with positive money value exists at all: whatever restraint an issuer announces, it will later prefer to abandon, and the market prices the abandonment from the start. Hellwig concludes that outside money is a natural monopoly in the loose sense that any other organisation of the market destroys the market itself, and suspects that time inconsistency is "the deepest and least solvable problem for the monetary constitution" (Hellwig, 1985, p. 582).

This is where the asymmetry between the two issuers of Stage 3 becomes decisive. A central bank operates under a politically imposed mandate, institutional rules, transparency requirements, and democratic oversight: its governance framework is precisely the binding precommitment device that Hellwig's analysis shows to be necessary for outside money to hold value at all. Its instrument q—equivalently, the credit expansion $\bar{D}$—is an expression of institutional commitment, not merely a policy choice. A private issuer elevated to outside-money status inherits the expansion incentive—each unit of outside money issued earns seigniorage at essentially zero marginal cost—but no comparable commitment technology. Nor is the redemption promise that defined the derivative stage a substitute once the escalation is underway: a guarantee disciplines only while the penalty for breaking it exceeds the gain, and Stage 2's dormant redemption right—rarely exercised, maintained for compliance—is exactly a guarantee whose breach grows cheap. We return to this enforcement margin in Section 6.

It is worth stating plainly what in this section is a result and what is not. The anchoring dichotomy of Section 4.1 and the indeterminacy (Proposition 3) are results: they characterise the boundary of the settlement hierarchy and the end-state beyond it. The transitions between the stages—redemption falling dormant, the alignment of issuer and fiscal incentives, the irreversibility of a built payment infrastructure—are institutional judgments on which the model is silent; so are the dynamics of a collapse, which the static frame cannot describe. The division is deliberate. What the model contributes is the location of the cliff, not a forecast of the fall: within the hierarchy, every effect of stablecoins is a bounded comparative static; beyond it, nominal determinacy itself is lost. Where the cliff lies is exactly the kind of question a model should be asked; when an economy reaches it is not.

5. Monetary policy within the anchored regime

Section 4 located the boundary at which the anchor itself is at stake. This section returns to the regime where the anchor holds: the nominally anchored system of Section 3, in which outside money remains the ultimate settlement medium and stablecoins move the terms of $M + qD + \theta B$.

What can the central bank do there with its own instruments? The monetary policy instrument is q – the price of raising liquidity from the banking system at $t = 0$ (or equivalently the volume of credit D). The monetary policy offset for a rise in the price level can be read from the accounting.

By Corollary 1 we have $\bar{D} = V + (1 - \theta)B$. The nominal value of trade is

$$V = \bar{D} - (1 - \theta)B$$

The offset can now be read off. Settlement absorption pins aggregate credit by the instrument alone, $\bar{D} = (M + B)/(1 - q)$, so at a given q a rise in the moneyness factor leaves $\bar{D}$ unchanged and raises nominal trade volume dollar-for-dollar with the bonds made spendable, $\Delta V = B\,\Delta\theta$. To hold V at its baseline value, credit must contract by exactly that amount, $\Delta\bar{D} = -B\,\Delta\theta$; since settlement absorption ties credit to the instrument, $\bar{D} = (M + B)/(1 - q)$, the instrument setting q' that implements the contraction solves $(M + B)/(1 - q') = \bar{D} - B\,\Delta\theta$ exactly, and to first order the required move is

$$dq \;=\; -\,(1 - q)\,\frac{B}{\bar{D}}\,d\theta.$$

The moneyness that enters this accounting is the effective θ_{eff} of Section 3.4: the offset the central bank must mount scales with the breadth adoption has actually reached, not with the engineered $\theta(S, \rho)$.

Two qualifications delimit what the offset achieves. First, the tier: as a statement about nominal volume the offset is distribution-free, resting only on Proposition 1(ii) and the absorption identity; reading it as a statement about the price level inherits the uniformity hypothesis of part (iii). Second, the contraction restores nominal volume exactly, but the price level only up to the real side: nominal volume is prices times real trade, and the wider wedge contracts real trade, so restoring the price level itself requires overshooting—the cost to which we now turn.

The offset is therefore feasible, and the framework locates its cost precisely. Three observations, in ascending order of importance. First, the central bank tightens against payment capacity it did not create: the additional liquidity enters through the coin's wrapping of the bond stock, outside the banking system, while the contraction falls on intermediated credit $\bar{D}$—on the trade that banks finance. Second, the offset accommodates a private expansion whose revenues it does not touch: the issuer continues to earn the yield on the wrapped bond stock, while the cost of the tightening is borne by bank-dependent trade. Third, and decisively, the instrument itself is a wedge: $r = (1 - q)/q$ separates buying prices from selling prices, and raising it to withdraw payment capacity suppresses exactly the activity the monetary system exists to enable. This is the offset-cost result of Dubey and Geanakoplos (2026): the price level can be defended with q, but only by forfeiting gains from trade. The summary of this section is accordingly not that policy is powerless but that its options carry a real cost—a cost the next section's discussion of policy postures takes as an input.

Table 5 makes both the offset and its cost computable in the running example, at the extreme $\theta = 1$ where the entire bond stock is spendable. The unoffset stock channel raises prices by the factor $\kappa(1) = 15/13$ with the allocation unchanged, as Proposition 1(iii) predicts. The volume-restoring instrument $q' = 7/13$, computed from the accounting condition above, returns nominal trade volume exactly to its baseline value—yet prices remain 4.5 percent above baseline: the wider wedge has contracted real trade, and the same nominal volume now finances fewer transactions. Restoring the price level itself requires overshooting to $q'' \approx 0.502$, an interest rate of 99 percent against the baseline's 67 and the accounting offset's 86, and the overshoot is paid in real terms: nominal volume falls 7 percent below baseline and both households lose utility (2.1 and 1.3 percent). The first and last rows of the table share a price level but not a real economy; the difference between them is the offset cost, measured. The equilibrium identities hold to solver precision throughout.

Table 5: The monetary policy offset and its cost in the running example ($M = 2.2$, $B = 1.1$, baseline $q = 0.6$), with the entire bond stock spendable ($\theta = 1$; equilibria computed via the equivalence of Proposition 1(i), dgme package v0.4.0). Row 2: the unoffset stock channel; prices rise by the factor $\kappa(1) = 15/13$. Row 3: the volume-restoring instrument $q' = 7/13$ from the accounting condition $(M + q'B)/(1 - q') + \theta B = (M + qB)/(1 - q)$, equivalently $\Delta\bar{D} = -\theta B$; nominal volume returns exactly to its baseline value while prices remain above baseline. Row 4: the price-restoring instrument q'', solved so that the geometric mean of the two price factors $p_\ell/p_\ell(\text{baseline})$ equals one. $r = (1 - q)/q$: the interest rate; index: geometric-mean price factor relative to the baseline row; V: nominal trade volume; $\bar{D}$: aggregate credit; u^1, u^2: household utilities; residual: largest absolute violation of the identities $(1 - q)V = M + qB + \theta(1 - q)B$, $\bar{D} = V + (1 - \theta)B$ and $(1 - q)\bar{D} = M + B$. Rows 1 and 4 have the same price level but not the same real economy: restoring prices costs both households utility — the offset cost made visible.

	q	r	p_1	p_2	index	V	$\bar{D}$	u^1	u^2	residual
baseline ($\theta = 0$)	0.60	0.67	2.61	2.18	1.00	7.15	8.25	2.51	1.62	$4.5e-12$
no offset	0.60	0.67	3.01	2.51	1.15	8.25	8.25	2.51	1.62	$5.3e-12$
volume offset q'	0.54	0.86	2.77	2.23	1.05	7.15	7.15	2.48	1.61	$2.5e-10$
price offset q''	0.50	0.99	2.68	2.12	1.00	6.63	6.63	2.46	1.60	$1.8e-09$

A further consideration is political rather than technical, and the model supplies its mechanics while remaining silent on its resolution. Corollary 2 showed that an established stablecoin sector makes new debt more potent per bond issued; read from the Treasury's side, the same fact is an enlarged and captive demand base for its paper. A central bank offsetting the stock channel therefore tightens against an expansion from which the fiscal authority benefits, and full offset presupposes a monetary authority willing to hold that line—an institutional judgment, not a result of the framework.

Finally, the instrument reaches consequences, not causes. Monetary policy can offset the capacity a low-ρ issuer adds; nothing in q bears on the issuer's choice of ρ itself, which responds to the yield on deployed reserves (Section 3.2). The two shut-off valves of Section 3.3—reserve composition and debt management—are regulatory and fiscal levers, not monetary ones. Nor does anything in q bear on the breadth of adoption: whether the acceptance leg gets built—the margin on which the effective moneyness θ_{eff} of Section 3.4 rests—belongs to regulation and infrastructure policy, the amplify decision of Section 6, and it is the cheapest prevention margin of the three, since declining to amplify requires no intervention at all. If the aim is to prevent the channels rather than absorb them, the relevant margins lie outside the central bank's instrument set.

The scope of all this should be stated as plainly as its cost. Everything in this section operates within the anchored regime: the channels of Section 3 move terms of $M+qD+\theta B$, and q can move them back, at a price. No setting of the instrument, however, reaches the settlement margin itself: whether ultimate settlement remains in public money is not a variable that q controls. Policy within the monetary order has a cost; policy about the monetary order is a different kind of question, and it is the subject of the next section.

6. Which monetary functions should be public?

The analysis of this paper allows us to discuss the policy questions raised by the emerging stablecoin sector in a principled way. One question can organise the discussion: in a monetary order that supports and enables economic activity, which monetary functions should be public, which can and should be private, and where does private initiative help rather than hurt? The model gives this question a precise location. The frontier between public and private runs between the payment constraint (T) and the settlement constraint (R): the channels of Section 3 operate inside (T) and stay bounded, while the escalation of Section 4 is precisely the migration of private instruments into (R), where nominal determinacy itself is at stake.

One remark on the status of what follows, before the discussion starts. The recommendations of this section are conditional in a precise sense: they combine the results of the model with a single objective—a price level that keeps its nominal anchor. Statements of the form "X must remain public" abbreviate the implication that, by the results of Section 4, no private organisation of X is compatible with that objective. Where an argument relies on an institutional judgment beyond the model, the text flags it as such. A reader who rejects the objective, or the model, is free to reject the conclusions—but will know exactly which premise carries them.

Payment innovations that mobilise future outside money for current payments are a recurring event, not an anomaly. Credit cards were such an innovation: they raise the credit volume D that a given base of outside money M can support (Dubey and Geanakoplos, 2010). Stablecoins are the most recent instance: they mobilise the deferred outside money embodied in government debt into payment capacity today. Empirically, stablecoins—despite their spectacular growth—are still used predominantly for transactions within the crypto universe and in the background processes of the payments industry; they have not yet reached day-to-day payments, and the problems identified in this paper are accordingly still mostly prospective. But such innovations do not come out of the blue. Successful ones cater to real needs, which is the ultimate reason for their adoption: this was the case for credit cards, and it may become the case for stablecoins once the appropriate infrastructure develops around them—the widely held expectation that stablecoins could lower the cost of cross-border payments is a case in point (Adrian et al., 2025; BIS, 2025).

Toward such innovations the public sector has, in principle, three postures. It can *suppress* them; it can *accommodate* them and offset effects such as the price-level push with monetary policy, at the cost established in Section 5; or it can *amplify* them through legislation, regulation and official endorsement—which for a payment instrument includes easing the build-out of the acceptance infrastructure on which its use in goods payments rests, the second leg of the two-sided market flagged at the opening of Section 3. The remainder of the section works out where each posture applies: why the settlement layer is a natural public monopoly and what kind of regulation follows on each side of that line (Section 6.1), the European dimension (Section 6.2), and central bank digital currency (Section 6.3).

6.1. The settlement layer as a natural public monopoly

Our analysis sheds light on where private initiative in payments can either be beneficial or at least a bounded, manageable externality. This is possible, in principle, whenever the initiative stays in the payment layer—in the model, an instrument that enters the payment constraint (T)—and leaves settlement in public money untouched.

The settlement layer and the unit of account, by contrast, are a natural monopoly in the loose sense of Hellwig (1985). As we argued in Section 4, *any* private organisation of the anchor necessarily destroys the anchor itself: two circulating outside monies leave the price level indeterminate (Proposition 3), a private issuer elevated to outside-money status inherits the expansion incentive without a commitment technology, and at a finite horizon there is no third option between redemption and a settlement role (the anchoring dichotomy of Section 4.1).

Regulation therefore faces two tasks of different kinds. In the payment layer, the task is ordinary conduct regulation of a legitimate business: rules on reserve composition ρ, on redemption timing, on disclosure. At the settlement layer itself, the task is not regulation but protection: settlement finality in public money, and taxes and bank obligations dischargeable only in the sovereign unit of account—commitments of a constitutional kind.

For stablecoins these two tasks translate into concrete rules, and the model says what each rule does. Reserve composition ρ decides what kind of thing the coin is. Backed by government debt, the coin remains a claim on sovereign money, and its value stands or falls with the sovereign anchor. Backed by assets that are nobody's nominal liability—gold, commodities, crypto-assets—it is a candidate money of its own: the competitor case of Section 4. Through $\theta(S, \rho)$, the same parameter also sets the size of the stock channel. Redemption terms decide whether the coin is money at all: a coin redeemable on demand at par can be accepted at par in payment—this is what moves θ off zero—while a coin redeemable only at settlement remains a mere claim. Disclosure rules support the par acceptance on which everything else rests. Read in this way, a regulatory rule answers a more precise question than "is this coin safe?": which term of $M + qD + \theta B$ does the rule move, and does it touch payment or settlement? Two examples display how the framework can help to sort out concrete policy issues.

First example: reserve backing. A coin fully backed by central bank reserves is a pass-through claim on M: no term of the payment capacity moves, the stock channel is switched off, and no marginal demand reaches the bill market, so the fiscal amplification of Corollary 2 is switched off with it—this is the neutrality valve $\rho = 1$ of Section 3.3, institutionally realised as a narrow bank, and it sits one step away from a central bank digital currency (Section 6.3). The choice between bill backing and reserve backing is therefore not a prudential detail: it is the switch on the inflation channel. It is worth noting that the financial-stability literature arrives at the same design from the opposite direction: in models of run risk, full reserve backing (with remunerated reserves to keep issuance viable) is the configuration that makes the coin run-proof (Li et al., 2026). The design that makes the coin safe from runs and the design that switches off its effect on the price level are one and the same, reached by two arguments that have nothing to do with each other. One qualification prevents a natural misreading: backing by central bank money is a statement about the composition of the issuer's assets—the parameter ρ—not about who holds an account at the central bank. The configuration can be implemented as a pass-through arrangement in which regulated custodians hold the reserves and the issuer holds claims on them (Adrian and Mancini Griffoli, 2019); whether issuers themselves receive direct access to the central bank's balance sheet is a separate access-policy decision on which the model is silent. An endorsement of reserve backing is therefore an endorsement of a monetary configuration, not of opening the central bank's balance sheet to stablecoin issuers.

A remark on the rails. Even the fully neutralised, reserve-backed coin remains a choice of payment *rails*: transfers run as tokens on blockchain infrastructure over the public internet. The benefits usually claimed for stablecoins—around-the-clock operation, programmability, cheap cross-border transfer—are benefits of these rails, not of the claim structure (Lagarde, 2026), and none of them requires the money to be privately issued or bond-backed. The rails raise a policy question of their own, separate from everything the model describes: a permissionless global chain works only as long as global connectivity does. Should international connectivity break apart, a domestic payment system backstopped by cash can fall back on itself; a payment system running on global blockchain infrastructure cannot. This is an operational judgment, not a monetary one, and we return to it briefly at the end of the section; here it serves as a caution that "monetarily neutral" does not mean "policy-irrelevant."

Second example: interest on coins. Should stablecoins be allowed to bear interest? Both GENIUS and MiCA say no: issuers must not remunerate holders. The accounting shows what is at stake in this choice. A holder of a bill-backed coin forgoes the yield on the very bills that back it; as long as this is so, holding coins is costly and demand for them stays limited. Allow the issuer to pass the bill yield through to holders, and the sacrifice disappears: the coin becomes as attractive as the bill itself, and demand—the scale S, and with it the moneyness $\theta(S, \rho)$—grows accordingly. In the limit the instrument is easy to describe: an interest-bearing, bill-backed coin is simply a government bond that circulates. This is the case $\theta \to 1$, and precisely the economy of Dubey and Geanakoplos' Theorem 4—the stock channel at its maximum. The interest ban, whatever its legal motivation, is therefore also a cap on the inflation channel: permitting pass-through of the bill yield would open the channel fully, banning it keeps the coin's scale, and with it the channel, small. Two limits of this reading should be stated. What coin interest would do to bank deposit funding belongs to the financial-stability bundle excluded from this paper. And how interest on a token would enter the one-period framework—as a coupon delivered at settlement, or as a discount in the coin's price—we have not modelled, so the statement concerns the direction of incentives, not an equilibrium result.

The existing frameworks. With the two examples in hand, we can return to the legislation actually on the books and read it in the same terms, neutrally: each provision is a rule on one of the margins just displayed—what backs the coin, and what its holder may earn. The U.S. GENIUS Act of 2025 consists of conduct rules throughout: licensing, one-to-one reserves in cash and short-term Treasuries, par redemption, monthly disclosure, and a prohibition on paying interest to holders. The purposes of these rules are plain and legitimate: a licensed perimeter for issuance, holders protected by full reserves, redemption at par and disclosure, and a legal foundation for payment innovation—joined, in the design's own strategic reading, by a strengthened international position of the dollar. The accounting adds the monetary side of the ledger, which this perspective leaves unpriced. Two observations follow. Its reserve menu steers issuers into the bill-backed design—the only permitted configuration that earns interest, and the one in which the stock channel operates and new debt arrives partly spendable—and its interest prohibition is the rule of the second example. What the Act does not address is the settlement margin: what may be discharged in coins. That silence is not necessarily a drafting flaw—the protections of the sovereign unit live in surrounding law, in tax law and legal tender statutes that predate the Act—but the analysis of this paper identifies exactly that margin as the crucial one, so a framework consisting entirely of conduct rules leaves the decisive protection implicit. The European regulation MiCA reads similarly in most respects—par redemption, reserve and disclosure requirements, an interest prohibition—with one provision of a different kind: caps on the use of tokens denominated in a foreign currency as a means of exchange.[27] This rule does not prescribe how a coin is to be run; it limits how far a foreign unit of account may spread. It is the one provision in either framework that guards the settlement side of the line, and it is motivated, tellingly, by monetary sovereignty rather than by investor protection.

This division of tasks is not only a deduction from the model; history ran the experiment. The national bank notes of the U.S. National Banking era (1863–1913) were private circulating liabilities backed by Treasury securities: bonds made spendable, in the language of Section 3, and the closest historical analogue of the bond-backed coin (Gorton and Zhang, 2023; Eichengreen, 2025). What the following century of institutional development settled is precisely the line drawn in this subsection: note issue migrated to the public sector, while innovation in the payment layer remained private. We do not claim that the historical motives were the model's—panics, the inelasticity of bond-backed issue, and seigniorage all played their part—but the division of labour that survived is the one the analysis identifies: private initiative in the payment layer, a public monopoly at the settlement layer.

6.2. The European dimension

In Europe the policy discussion on stablecoins has recently been gaining momentum. That debate is usually framed around a striking difference in scale: euro-denominated coins stood at roughly 450 million euro of market capitalisation in early 2026, against some 300 billion dollars for their dollar counterparts—a gap of several hundred times (Born et al.,

[27] Chapter and verse, from Regulation (EU) 2023/1114 (MiCA): the interest prohibitions are Articles 40 (asset-referenced tokens) and 50 (e-money tokens), both extending to crypto-asset service providers and defining interest broadly. The means-of-exchange caps are Article 23(1): once the use of a token as a means of exchange within a single currency area exceeds, as a quarterly average, one million transactions and EUR 200 million in value per day, the issuer must stop issuing and submit a plan to bring usage back below the thresholds; Article 58(3) applies these caps to e-money tokens denominated in a currency that is not an official currency of a Member State—the dollar coin in Europe, not the euro coin. Article 24(3) requires authorities to cap issuance outright once the central bank finds a threat to payment systems, monetary policy transmission, or monetary sovereignty.

2026). Some read this gap as a call to catch up (Schaaf, 2025; Bindseil and Voloder, 2026); others see euro stablecoins less as money than as a vehicle for renewing the technology of wholesale payments and securities settlement, and for strengthening the international role of the euro.

The gap in scale has a second reading, and for European policy makers it is the more immediate one: nothing confines the dollar coins to the dollar economy—in emerging markets, stablecoin inflows already track the drivers of classic deposit dollarisation while slipping past capital controls (Hofmann et al., 2026b). Should they migrate into ordinary euro-area payments, the event would be of a different kind from anything in Section 3. A dollar coin at a euro-area checkout moves no term of the euro's payment capacity $M + qD + \theta B$—its backing is US bills, and the capacity it expands is the dollar's. It is invisible in the euro accounting, and consequential for exactly that reason: what it erodes is not the euro's price level but the domain of trades the euro mediates—currency substitution, a competing-money event of the family of Section 4. And it enters the escalation ladder with an advantage no private candidate money has: it need not survive the loss of a redemption anchor, because it imports an anchor another sovereign already maintains, so the discipline problem of Section 4.3 does not arise for it. The protections that section identifies—conversion frictions and mandatory settlement in the sovereign unit—are precisely what MiCA's means-of-exchange caps on foreign-denominated tokens (Section 6.1) write into law. Read through the model, those caps are not a market-share technicality but the settlement-side guard against a competitor that arrives pre-anchored. As throughout this section, this is a reading of the model's structure, not a result: the model describes a single economy, not two currency areas.

A clear perspective on these questions matters for Europe for strategic reasons. The United States has placed active political and regulatory support behind its stablecoin industry, reinforced by the GENIUS Act and by the expectation that dollar-backed coins will entrench the dollar's international position. In terms of the three postures, this policy is amplification: it steers issuers into the bill-backed design in which new public debt arrives partly spendable (Corollary 2), so the hoped-for strengthening of the currency may prove self-defeating—a conclusion Angeloni (2026) reaches by a different route, reading the same policy as part of a drift towards fiscal dominance that would erode rather than entrench the dollar's international status. If so, the catching-up reading of the European gap loses its point. The question that remains is what a euro-denominated coin would be in the monetary structure the union actually has — and there the European case differs from the American one not in scale but in the architecture of its settlement layer.

The monetary union in Europe—unlike the United States—has a single settlement layer through the euro and the European Central Bank. In the language of our model, Europe has one (R). But unlike the United States it has no common fiscal policy: the

European settlement layer rests on twenty national debt stocks B_i with no federal B. This has two immediate consequences. The model contains a single stock of public debt and is silent on a union of several; the consequences are readings of its structure, not results.

The first is that a euro stablecoin would have to pick its backing from heterogeneous national debts. If coin demand concentrates on the scarcest safe asset—German Bunds, say—funding costs are redistributed among the members of the currency union: a flow channel with no counterpart in the United States.

The second consequence would be riskier still. Should a European stablecoin market gain scale, the same financial engineering would allow national currencies to reappear in the guise of payment technology. Consider a coin backed by a single member's debt—a Deutschmark coin built on Bunds—and accepted in settlement within one region. Such a coin would re-establish exchange rates inside the union and, carried to its end, unravel the common currency. In the model's terms, it would split the single (R) along national lines.

This second consequence is not a theorist's invention; the euro area has already seen a small-scale rehearsal, before any blockchain was involved. In May 2019 the Italian Camera dei Deputati approved, with the support of government and opposition alike, a motion—an act of political direction, not legislation—on the settlement of the public administration's commercial arrears, which listed, among the instruments to be evaluated, small-denomination, non-interest-bearing government securities, printed like banknotes and receivable at par for tax payments—known in the public debate as mini-BOTs.[28] Nothing was ever issued, and the motion asked only for an evaluation. But the design—euro-denominated state paper meant to circulate—was immediately read as the first step towards a parallel national currency and, in the limit, towards an exit from the euro (Cecchetti and Schoenholtz, 2019), and the proposal was dead within two weeks, under the verdict of the finance minister, the Banca d'Italia and the ECB. Mario Draghi's formulation of that verdict—at the time of this episode he was president of the ECB—is the anchoring dichotomy of Section 4.1, stated from the podium: mini-BOTs "are either money and then they are illegal, or they are debt and then that stock goes up" (Draghi, 2019). The episode teaches two things. The incentive to resurrect a national unit inside the union is political, and does not wait for technology. And the institutional defence held because the instrument was recognised for what it was; a stablecoin version of the same design would be technically easier to launch and, dressed as payment innovation, harder to recognise.

The regulatory environment adds a further European particularity. MiCA does not only differ from the GENIUS Act in its unit-of-account caps (Section 6.1); it also mandates

[28] The parliamentary act is Mozione 1-00013 (Baldelli and others), approved by the Camera dei Deputati on 28 May 2019. We do not cite the Italian parliamentary record directly and rely on English-language accounts of the episode, in particular Cecchetti and Schoenholtz (2019) and the ECB press conference transcript cited below.

a different composition of the backing. An issuer of a euro e-money token must hold at least 30 percent of the funds received as deposits with credit institutions—60 percent for significant tokens—and may invest only the remainder in high-quality liquid assets such as government bills.[29] Read on the reserve-composition margin ρ of Section 6.1, this is a constraint with monetary content. The deposit share of the backing is a claim on inside money: it repackages existing D and mobilises no future outside money, so the stock channel operates only on the bill share. Where the GENIUS reserve menu steers issuers into the configuration in which the channel is fully open, MiCA caps the bill-backed share of issuance at seventy percent—forty for significant coins—by construction. The motivation is prudential, deposits being the liquid tranche for redemptions, not monetary; but as with the run-proofing convergence of Section 6.1, a rule written for stability has a side effect on the inflation channel, this time a throttling one. The necessary caveat: the deposit share imports bank credit risk into the coin's backing, a matter that belongs to the financial-stability bundle this paper excludes.

What is actually emerging in Europe fits this reading. In September 2025 a consortium of European banks—initially ING, UniCredit, CaixaBank, Danske Bank, SEB, KBC, DekaBank, Banca Sella, and Raiffeisen Bank International, joined since by BNP Paribas, BBVA and others, thirty-seven banks from fifteen countries by May 2026—formed a joint venture, Qivalis, domiciled in the Netherlands, to issue a MiCA-compliant euro stablecoin, with first issuance targeted for the second half of 2026 (Caixabank, 2025). The issuers are regulated banks that already create inside money, and this changes what the coin is in the accounting. Backed by assets the banks already hold—and, for institutions with access to the deposit facility, ultimately by central bank reserves—the token is closer to a re-denomination of existing claims onto new rails than to a new monetary claim: the configuration sits at or near the neutrality valve $\rho \approx 1$ of Section 6.1, and no term of $M + qD + \theta B$ need move. The design is nonetheless privately attractive under MiCA's interest ban: the backing earns the deposit facility rate while token holders receive no remuneration, so the interest on the backing accrues to the issuers—the same margin that funds the US issuers, available here without activating the stock channel. The initiative thus combines monetary neutrality with a viable business model, and its stated motivation—strategic autonomy in payments, reducing dependence on non-European card networks and dollar-denominated infrastructure—is a question about the rails and the governance of the payment layer, the subject of the rails remark in Section 6.1, not

[29]Article 54 of MiCA: at least 30 percent of the funds received always on deposit in separate accounts with credit institutions, the remainder in secure, low-risk, highly liquid financial instruments denominated in the currency the token references. For significant e-money tokens the deposit floor rises to 60 percent, via Article 58(1), point (a), in conjunction with Article 45(4) and the technical standards mandated by Article 45(7), first subparagraph, point (b). Both floors sit in the Regulation itself; the regulatory technical standards operationalising the reserve liquidity requirements were still pending in mid-2026.

about the anchor. One differentiation matters for what follows: the venture's payment focus is cross-border and business-to-business settlement and digital-asset markets, not the retail checkout. The private venture the retail digital euro actually faces is the European banking sector's wallet alliance EPI/Wero, the most prominent bank-backed retail payment initiative at the European level, while Qivalis operates in the territory of the Eurosystem's own wholesale DLT initiatives, Pontes and Appia (Cipollone, 2026). Each private venture, in other words, has a public counterpart on its own margin—built, or being built.

The two consequences of the missing federal debt flow from a single structural fact: any bond-backed euro coin must pick its fiscal anchor from among the members, and every pick has distributional or, in the limit, secessionist content. This is what places the digital euro in the present argument. As a liability of the union's central bank it is the one euro-denominated token exempt from that choice: not a θ-instrument on any member's debt but a new form of M itself, issued at the settlement layer the union already shares. It therefore raises neither of the two consequences—it selects no national debt, and it can resurrect no national unit—while offering the modern rails that motivate the second camp in the European debate. The digital euro is, of course, an instance of a central bank digital currency, and what the model says about this class of instruments in general is the subject of the next subsection.

6.3. Central bank digital currency

A central bank digital currency (CBDC) usable in retail payments, such as the digital euro, is a central bank liability held directly by households. In terms of our model it is an institutional implementation of the compression of Section 2: outside money M held directly by households. Consequently there can be no price-level effect: no term of $M + qD + \theta B$ is moved by the introduction of a CBDC, which changes the form of M, not its size—the payment-capacity counterpart of the equivalence and irrelevance results of Brunnermeier and Niepelt (2019) and Niepelt (2026). The effects a CBDC does have—most prominently the possible migration of bank deposits into central bank money—live inside the two-tier structure the model compresses away, are the subject of the CBDC design literature (Andolfatto, 2021; Fernández-Villaverde et al., 2021; Keister and Sanches, 2023), and belong to the financial-stability bundle this paper excludes. What remains is not a monetary event but a hierarchy event: a public option entering the payment layer described by the payment constraint (T).

The instruments that have appeared in this paper can be ordered by two questions about any circulating token: who issues it, and on whose anchor does its value rest? Crossing the two questions yields Table 6. Sovereign money occupies the public–public cell: cash and central bank reserves today, joined by a retail CBDC and, at the wholesale level, by a tokenised cash leg of the kind the Eurosystem is building (Section 6.2)—new entries in the cell, not new cells. The regulated stablecoin of this paper is a private instrument on the public anchor: the *derivative* of Section 4, and the cell it shares with tokenised deposits and with inside money generally, whose value likewise rests on redemption into the public anchor. The Stage-3 coin is the private–private cell: a token circulating as outside money in its own right, the *competitor* against which Proposition 3 and the anchoring dichotomy of Section 4.1 warn. The fourth cell—a public instrument on a private anchor—is essentially empty, and its emptiness is instructive: a sovereign issuer that borrows an anchor not its own has already given up what the protections of the settlement layer exist to keep, and the nearest real configurations, currency boards and dollarisation, are public instruments on a *foreign public* anchor, not on a private one.

	Anchor	
Instrument	Public	Private
Public	sovereign money cash, reserves, retail CBDC, wholesale cash token	(essentially empty)
Private	stablecoin: *derivative* bond- or reserve-backed coin, tokenised deposits, inside money	Stage-3 coin: *competitor* real-asset-backed or unbacked circulating token

Table 6: Payment instruments classified by the issuer of the instrument and the anchor on which the instrument's value rests. Escalation (Section 4) is movement along the anchor axis; a CBDC is movement along the instrument axis.

The point of the table is that it makes directions of movement visible. The escalation path of Section 4 is movement along the *anchor* axis with the instrument held private: from derivative to competitor, the passage whose price the indeterminacy proposition and the anchoring dichotomy state. A CBDC is the orthogonal move: along the *instrument* axis, with the anchor held fixed. This restates the paragraph that opened this subsection: on this path nothing monetary happens—no term of $M + qD + \theta B$ moves—and what changes is who operates in the payment layer. The table also locates the fully reserve-backed coin of Section 6.1: at $\rho = 1$ the private instrument is bound so tightly to the public anchor that it sits on the boundary of the sovereign cell—the configuration the literature calls a *synthetic* CBDC (Adrian and Mancini Griffoli, 2019)—separated from the public instrument by legal form, not by monetary substance.

The public debate on CBDCs and stablecoins often runs on a familiar contrast: bureaucratic, overregulated Europe enters the digital age with a public-sector product nobody has asked for, while the innovation-friendly United States chooses stablecoins, the technology the market actually wants. This framing treats the CBDC and the stablecoin as substitutes—two rivals for a single role, differing along the instrument axis of Table 6. What the analysis shows is that they are moves along different axes. A CBDC leaves the price level untouched but changes the hierarchy: it places a new public instrument in the payment layer. A redeemable stablecoin does exactly the reverse: it can move the price level through the channels of Section 3, while leaving the hierarchy intact for as long as the redemption promise is kept. The two instruments answer different questions and can coexist; the substitution framing is a category error.

One part of the contrast survives this correction: the objection that nobody has asked for the public product. It deserves a direct answer rather than counter-advocacy. The demand the market reveals for stablecoins is largely a demand for rails (Section 6.1), and rails are available on any anchor; what revealed demand does not price are the costs attached to the claim structure—the stock channel, the fiscal amplification, the hierarchy risk—which fall on the holders of the currency at large, not on the buyers of the coin. Revealed demand therefore cannot settle the policy question. And the concession runs the other way too: the model supplies no need-based case for a CBDC either. Its rationale is not that anybody asked for it; it is insurance—keeping a public instrument present in the payment layer. Whether a digital euro would be a well-executed product is a question about rails and institutional delivery on which our framework is silent.

A second position is content to wait and see on the retail CBDC but holds that wholesale finance, at least, needs stablecoins: tokenised markets require a tokenised settlement asset, and the coins are the candidate at hand. Our analysis turns this position exactly on its head. Wholesale settlement is not a peripheral technical domain in which private experimentation is cheap; it *is* the settlement layer. A private settlement asset accepted there moves θ at the very margin where finality lives—the most corrosive place it can move. Retail payments, by contrast, belong to the payment layer, where private initiative is the legitimate, conduct-regulated business of this section. The position thus recommends private money precisely where the analysis reserves the public monopoly, and tolerates the public instrument only where private initiative was harmless to begin with. It has, moreover, been overtaken by events: a tokenised settlement asset does not require a private issuer, and the central-bank-money rail for DLT settlement is being built—the Eurosystem's Pontes infrastructure, with finality in central bank money through the TARGET services (Section 6.2)—placing a public cash token in the public–public cell of Table 6 at the wholesale level.

These wholesale initiatives and the retail digital euro are parts of one strategy—keeping central bank money fit for a digital world—not rival tracks (Cipollone, 2026), and together they work like insurance: they keep a public instrument available at both levels of the payment system, retail and wholesale. Should private rails come to dominate retail payments—the banking sector's wallet alliance, tokenised deposits, regulated coins—the public infrastructure preserves the ability of the sovereign to re-enter the payment layer at either level without rebuilding it from scratch. Like all insurance, the option has value precisely in the states of the world its critics consider unlikely.

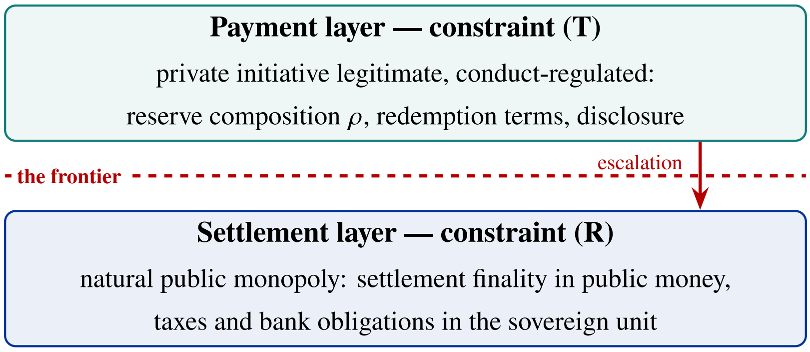

Figure 6: The frontier. The payment layer, governed by the transactions constraint (T), is where private initiative is legitimate and conduct regulation operates; the settlement layer, governed by the settlement constraint (R), is a natural public monopoly (Section 6.1). The red arrow is the migration of a private instrument across the frontier—the escalation of Section 4. The figure locates the first two clauses of the paper's stance: regulate conduct above the line, protect below it. The third clause, do not amplify, is a constraint on behaviour within the payment layer, not a place on the map.

The stance of this paper can now be stated in three clauses, each read off the accounting and the table; Figure 6 places the first two on the map. In the payment layer, regulate conduct: reserve composition, redemption terms, disclosure—rules that move ρ, θ and S—for a redeemable coin is a legitimate business of the same family as inside money. At the settlement layer, protect: settlement finality in public money, taxes and bank obligations dischargeable only in the sovereign unit—for this is the margin on which, by Proposition 3 and the anchoring dichotomy of Section 4.1, the task is not to regulate a market but to preserve an institution. And do not amplify: of the three postures, amplification is the one for which the analysis provides no support—legislation and official endorsement that boost demand for the bond-backed configuration, in which new public debt arrives partly spendable (Corollary 2). The question is meanwhile no longer an American one alone: coin projects denominated in other currencies are under way—in April 2026 a group of Swiss banks, UBS and PostFinance among them, opened a joint sandbox for a Swiss franc coin (UBS Switzerland, 2026)—while actual use of coins in ordinary payments remains, as noted at the outset of this section, negligible. The problem is largely prospective, and that is not a reason for comfort but the reason to act at the constitutional margin now: lines are easiest drawn while no one is standing on them. There remains, finally, the caution promised in Section 6.1: all of these instruments, public and private alike, run on rails, and how resilient those rails remain if global connectivity breaks apart is an open operational question—a judgment outside our model, but one a policy for the settlement layer cannot delegate. A domestic payment system backstopped by public money and public rails can fall back on itself; a payment system woven into permissionless global infrastructure cannot. Keeping the anchor public is, in the end, also the precondition for keeping that choice open.

7. Conclusion

Do stablecoins affect the price level? The answer this paper gives is yes. The channels that create this effect are all encoded in one equation describing the economy's payment capacity: $M + qD + \theta B$, the stock of outside money, commercial bank credit and the spendable part of government bonds. Two things must come together to make government bonds spendable. The financial engineering of stablecoins supplies the instrument; policy and regulation decide whether an acceptance infrastructure is built around it. The parameter θ_{eff} records how far both have gone. The stock channel makes part of the existing bond stock spendable, and the flow channel makes new debt arrive spendable (Prop. 1, Cor. 2). Adoption decides how much of this engineered spendability becomes economically effective.

All these channels are bounded effects within a nominally anchored monetary order in which the central bank is able to influence the price level through monetary policy setting interest rates. The escalation of Section 4 is not a larger effect but an effect of a different category. This escalation remains a risk as long as no rules are in place to prevent a migration of stablecoins from the payment layer (T) into the settlement layer (R) itself. When this migration happens the monetary order loses its nominal anchor and the price level becomes indeterminate.

What can policy do to deal with the growing stablecoin sector? Within the current monetary order monetary policy can in principle always act to offset the effects of inflationary finance. The costs of this counteracting policy are also clear: monetary policy needs to lean on an instrument that defends the price level but also forfeits gains from trade (Section 5). The policy guardrails preventing a migration of stablecoins from the payment layer into the settlement layer are constitutional in kind: if the price level is to keep its anchor, the rules of the monetary order must preserve the settlement layer for public money, while they can allow private initiative in the payment layer (Section 6).

None of this is a forecast. It is rather a map helping to navigate the policy landscape of stablecoins. The inflationary effects are dormant while coins stay in their crypto niche. How far stablecoins migrate into ordinary retail payments is shaped by the regulatory frameworks now taking shape, MiCA in the EU and the GENIUS Act in the US: policy can enable this migration, make it difficult, or rule it out. Adoption can also arrive without any official push, through merchants' own acceptance decisions; the choice policy cannot avoid is then whether to permit, restrain or forbid what emerges—a decision by default rather than by design. Adoption is in this sense a policy choice, and the paper is a decision aid for that choice. The constitutional lines that protect the nominal anchor and the monetary order itself are best drawn before anyone stands on them.

The Dubey–Geanakoplos model – on which our analysis is based – is a compressed one-period model which is able to reveal the decisive points most clearly. Its dynamics are only imagined, compressed within the single period (see the paragraph on the model's imagined dynamics in Section 2). It is no welfare analysis of stablecoins as such. What the model delivers is the payment capacity accounting that makes the questions raised analytically tractable. The framework and its cornerstone results are Dubey and Geanakoplos (2026)'s; what this paper adds is built inside their model: the two-date translation of the stock channel with its wealth-effect qualification (Proposition 1), the $\theta(S, \rho)$ endogenisation, the flow channel, the two-money indeterminacy result (Proposition 3), and the policy mapping. For the questions raised at the outset – who provides the nominal anchor, and what monetary order should support our economy – the model's final lesson is this: a monetary order is not found, it is chosen.

A. Proofs

Lemma 1 (No idle money)**.** *Fix prices (p, q) with $p \in \mathbb{R}^L_{++}$ and $q \in (0, 1)$. If a plan (x^h, ℓ^h, β^h) satisfies (T), (R) and (F), then there exists a plan $(x^h, \tilde{\ell}^h, \tilde{\beta}^h)$ with the same consumption bundle that satisfies (T), (F) and the tight settlement constraint (R′). The projection of the budget set onto consumption bundles is therefore unchanged when the unspent-money term μ^h is dropped from (R).*

Proof. Both (T) and (R) depend on (ℓ^h, β^h) only through the net financial position $\psi^h \equiv \ell^h + \beta^h$: constraint (T) reads $p \cdot \tau^{h,+} \le m^h - q\psi^h$, and (R) reads $p \cdot \tau^{h,-} + b^h + \psi^h + \mu^h \ge 0$ with $\mu^h = m^h - q\psi^h - p \cdot \tau^{h,+}$.

If $\psi^h \ge 0$, the tight constraint (R′) holds trivially, since all its terms are non-negative. Suppose $\psi^h < 0$, and set $\delta \equiv \min\{-\psi^h,\ \mu^h/q\}$ and $\tilde{\psi}^h \equiv \psi^h + \delta$. The new plan satisfies (T), since unspent money falls to $\tilde{\mu}^h = m^h - q\tilde{\psi}^h - p \cdot \tau^{h,+} = \mu^h - q\delta \ge 0$; and the left-hand side of (R) changes by $\delta - q\delta = (1 - q)\delta \ge 0$, so (R) continues to hold. If $\delta = -\psi^h$, then $\tilde{\psi}^h = 0$ and (R′) holds trivially. If $\delta = \mu^h/q$, then $\tilde{\mu}^h = 0$, so (R) and (R′) coincide, and (R′) holds. In either case, choose the decomposition $\tilde{\ell}^h = \min\{\tilde{\psi}^h, 0\} \le 0$ and $\tilde{\beta}^h = \max\{\tilde{\psi}^h, 0\} \ge 0 \ge -b^h$, which satisfies (F). The consumption bundle x^h is unchanged.

Conversely, any plan satisfying (R′) satisfies (R), since $\mu^h \ge 0$. The two budget sets therefore have the same projection onto consumption bundles. □

Remark 1 (Specialization of the Dubey–Geanakoplos economy)**.** *The economy of Section 2 is an instance of the economy of Dubey and Geanakoplos (2026), obtained by parameter settings their framework explicitly admits, together with a relabelling of the timeline.*

(i) Parameter settings. *Credit card limits $\zeta^h = 0$ for all h: credit cards are absent (their §2.5 allows exogenous limits; setting them to zero removes the instrument). Setup costs $\kappa^h = 0$: borrowing carries no shoe-leather cost (their §2.9). Loan limits $\chi^h = \infty$ for all h: bank credit is unrestricted, which they explicitly allow and single out as crucial for existence (their §2.3). Consequently the set $U = \{h : \chi^h = \infty\}$ of unrestricted agents is all of $\mathcal{H}$, and the ownership condition of their existence theorem—every good owned by some agent in U—reduces to our assumption $\sum_h e^h \gg 0$.*

(ii) Market structure. *Of their instrument set, the markets open here are the money–goods markets for every good, the loan market, and the bond–money market. Markets for direct purchases of goods with bonds are closed in the baseline; barter markets are closed throughout, as they impose.*

(iii) Timeline. *Their five intervals map onto the two dates as follows: intervals 1 and 2 (borrowing; trade) collapse into $t = 0$; interval 3 (credit card settlement) is empty since $\zeta^h = 0$; interval 4 (bank settlement, with netting of government bond deliveries) is $t = 1$;*

consumption (interval 5) follows settlement. The netting of bond deliveries against bank debt in their interval 4 is preserved in constraint (R). Their cash-in-advance constraint (2a)—purchases limited to money on hand—is (T); their settlement constraint (4) is (R); the carried balances that link their intervals reduce to the single term μ^h, which Lemma 1 eliminates.

(iv) Agent space. *They work with a continuum of copies of finitely many agent types (their §2.1) because the setup cost κ^h introduces a nonconvexity whose resolution requires within-type mixing in equilibrium (their §3.4). With $\kappa^h = 0$ the household problem is convex, type-symmetric equilibria exist, and a type-symmetric equilibrium of the continuum economy induces, type by type, a price-taking equilibrium of the finite economy with the same prices and allocations, and conversely.*

Results of Dubey and Geanakoplos (2026) cited in the text therefore apply to the present economy; in particular, the existence theorem is invoked on this basis. The role of this remark is to license those citations and to make the relationship to the published framework precise; the derivations specific to this paper's extensions are given directly in the text and in this appendix.

Proof of Proposition 1. Throughout, fix prices (p, q) with $p \in \mathbb{R}^L_{++}$ and $q \in (0, 1)$, and write $W^h(s) \equiv m^h + qb^h + (1 - q)s$ for $s \in [0, \theta b^h]$.

Step 1: household equivalence at fixed coin use. Substituting $(\tilde{m}^h, \tilde{b}^h) \equiv (m^h + s^h,\ b^h - s^h)$ transforms (T_θ) into (T), (R'_θ) into (R′), and the bound $\beta^h \geq -(b^h - s^h)$ into $\beta^h \geq -\tilde{b}^h$. Hence, for fixed s^h, household h's constraint set in $\mathcal{E}(\theta)$ is exactly the baseline budget set (1) of a household endowed with $(\tilde{m}^h, \tilde{b}^h)$. In particular, Lemma 1 applies verbatim, with $(\tilde{m}^h, \tilde{b}^h)$ in place of (m^h, b^h).

Step 2: collapse and full coin use. By Step 1 and the collapsed representation of Section 2 applied at the endowments $(\tilde{m}^h, \tilde{b}^h)$, a consumption bundle is attainable at coin use s^h if and only if

$$p \cdot \tau^{h,+} - q\, p \cdot \tau^{h,-} \ \leq\ \tilde{m}^h + q\tilde{b}^h \ =\ W^h(s^h).$$

Since $1 - q > 0$, W^h is strictly increasing in s^h, so these budget projections are nested and largest at $s^h = \theta b^h$. A bundle is therefore attainable in $\mathcal{E}(\theta)$ if and only if it satisfies the constraint with intercept $W^h(\theta b^h) = m^h + qb^h + \theta(1 - q)b^h$, and with strictly increasing preferences we may take $s^h = \theta b^h$ at any optimum. This is the liquid-wealth claim in (ii).

Step 3: part (i). By Steps 1–2, household h's attainable consumption bundles and optimal choices in $\mathcal{E}(\theta)$ coincide with those of the baseline household endowed with $(m^h + \theta b^h,\ (1 - \theta)b^h)$. Goods-market clearing is a condition on consumptions alone. On the financial side, only the net position $\psi^h = \ell^h + \beta^h$ enters the household's constraints, and at $s^h = \theta b^h$ the feasible decompositions (ℓ^h, β^h) coincide in the two economies (the bound is $\beta^h \geq -(1-\theta)b^h$ in both). Any decomposition satisfying bond-market clearing $\sum_h \beta^h = 0$ and supporting aggregate credit $\bar{D} = -\sum_h \psi^h$ in one economy therefore does so in the other, and the accommodation condition of the banking sector is identical. The two economies have the same monetary equilibria.

Step 4: part (ii). At an optimum the collapsed constraint binds: were it slack, strictly increasing preferences would afford a better bundle. Summing the binding constraints over households and using $\sum_h p \cdot \tau^{h,+} = \sum_h p \cdot \tau^{h,-} = V$ gives

$$(1-q)\,V \;=\; \sum_h W^h(\theta b^h) \;=\; M + qB + \theta(1-q)B,$$

which is strictly increasing in θ whenever $B > 0$.

Step 5: part (iii). Uniform composition means $b^h = c\,(m^h + qb^h)$ for all h with a common $c \geq 0$; summing over households gives $B = c\,(M + qB)$, so $c = B/(M + qB)$ and $\kappa(\theta) = 1 + \theta(1-q)c$. Let $(\bar{p}, q)$ with allocation $(\bar{x}^h)_h$ be a monetary equilibrium of the baseline at the policy rate q. Consider $\mathcal{E}(\theta)$ at prices $(\kappa\bar{p}, q)$ with $s^h = \theta b^h$. For each h, the collapsed constraint reads

$$\kappa\bar{p} \cdot \tau^{h,+} - q\kappa\bar{p} \cdot \tau^{h,-} \;\leq\; W^h(\theta b^h) \;=\; \big(1 + \theta(1-q)c\big)(m^h + qb^h) \;=\; \kappa\,(m^h + qb^h);$$

dividing by $\kappa > 0$ returns the baseline collapsed constraint at prices $\bar{p}$. The set of attainable net trades is unchanged, so $\bar{x}^h$ remains optimal for every h, and goods markets clear. Financial positions supporting the allocation exist as in Step 3, scaled to the new nominal volume, and the banking sector accommodates the implied aggregate credit at the unchanged policy rate q. Hence $(\kappa\bar{p}, q)$ with allocation $(\bar{x}^h)_h$ extends to a monetary equilibrium of $\mathcal{E}(\theta)$. Finally, $\kappa(\theta) > 1$ if and only if $\theta > 0$ and $B > 0$. □

Proof of Corollary 1. As in the settlement-absorption argument of Section 2, both slacks of the constraint system vanish at an optimum: unspent liquidity is zero because idle money could buy more, and surplus settlement capacity is zero because resources beyond obligations are worthless at the horizon. Sum the two zero-slack conditions across households, using goods- and bond-market clearing and $\sum_h s^h = \theta B$ (Step 2 of the preceding proof).

At $t = 0$, aggregate purchases V decompose into coin-financed purchases θB and money-financed purchases, and all money on hand is spent:

$$V - \theta B \;=\; M + q\bar{D},$$

where bond trades net out by $\sum_h \beta^h = 0$. At $t = 1$, all settlement resources are pledged to the banking sector: sales proceeds V—paid in money or in coins, each delivering one unit of outside money at settlement—plus the deliveries $(1-\theta)B$ on retained bonds equal the face value of bank debt:

$$\bar{D} \;=\; V + (1-\theta)B.$$

Combining the two,

$$(1-q)\bar{D} \;=\; \big[V + (1-\theta)B\big] - \big[V - \theta B - M\big] \;=\; M + B.$$

The settlement absorption condition (6) is unchanged: the coin alters the composition of payments and, through Proposition 1, the nominal volume of trade, but the interest on inside money still absorbs the economy's entire outside wealth at the horizon. □

Proof of Proposition 2. Fix prices (p, q) and a household h with placements (d^h, w^h), $d^h \in [0, m^h]$, $w^h \in [0, \theta b^h]$, $\theta \equiv (1-\rho)S/B$.

Step 1: the deposit cancels. At $t = 0$ the household's means of payment are its remaining money $m^h - d^h$, its borrowing proceeds $-q\ell^h$, and its tokens $d^h + w^h$ at par, so the payment constraint reads

$$p \cdot \tau^{h,+} + q\beta^h \;\leq\; (m^h - d^h) - q\ell^h + (d^h + w^h) \;=\; m^h - q\ell^h + w^h,$$

which is (T_θ) with $s^h = w^h$. At $t = 1$ a token, whether held unspent or received as sales proceeds, is a claim to one unit of outside money on the issuer, and the issuer's resources—the impounded money ρS plus the deliveries $(1-\rho)S$ on the wrapped bonds—equal the tokens outstanding S, so every such claim is discharged at par. The household's settlement resources are therefore its sales proceeds (money and tokens alike delivering at par), the deliveries $b^h - w^h + \beta^h$ on its retained bond position, and its unspent balances; against them stands the repayment $-\ell^h$. Dropping the unspent balances by the argument of Lemma 1 yields (R'_θ) with $s^h = w^h$, and the no-short-sales bound is $\beta^h \geq -(b^h - w^h)$ as there. The constraint system thus depends on (d^h, w^h) only through w^h, which is part (i); in particular, any deposit profile with $\sum_h d^h = \rho S$ (one exists since $\rho S \leq M$) supports the same household choices.

Step 2: reduction to the coin economy. By Step 1, the issuer economy at the proportional cap $w^h \leq \theta b^h$ has exactly the constraint system of $\mathcal{E}(\theta)$, so the two have the same monetary equilibria. By Step 2 of the proof of Proposition 1, $w^h = \theta b^h$ at any optimum; summing, $\sum_h w^h = \theta B = (1-\rho)S$, so the issuer's balance sheet is reproduced in equilibrium and the cap is consistent with the aggregate placements. The condition $(1-\rho)S \leq B$ gives $\theta \in [0, 1]$. The identity of part (ii) is Proposition 1(ii) at $\theta = (1-\rho)S/B$, using $\theta(1-q)B = (1-q)(1-\rho)S$; the price statement is Proposition 1(iii) at the same value; and the absorption condition transfers by Corollary 1. □

Proof of Corollary 2. Part (i). The identity $(1-q)V = M + qB + \theta(1-q)B$ is Proposition 1(ii) at the stock B; the remaining statements are its derivatives. The marginal contribution of a unit of bond face value is $\partial\big[(1-q)V\big]/\partial B = q + \theta(1-q)$; its ratio to the value at $\theta = 0$ is $\big(q + \theta(1-q)\big)/q = 1 + \theta(1-q)/q$; and the cross effect is $\partial^2\big[(1-q)V\big]/\partial B\,\partial\theta = 1 - q > 0$.

Part (ii). Uniform composition at the stock B means $b^h = c\,(m^h + qb^h)$ for a common c, equivalently $b^h = \lambda m^h$ with $\lambda = c/(1-qc)$; uniformity after issuance likewise gives $\tilde{b}^h = \tilde{\lambda} m^h$ for a common $\tilde{\lambda}$. The effective liquid wealth of Step 2 of the proof of Proposition 1 is then proportional to m^h at both stocks,

$$W^h \;=\; m^h\big(1 + (q + \theta(1-q))\lambda\big), \qquad \tilde{W}^h \;=\; m^h\big(1 + (q + \theta(1-q))\tilde{\lambda}\big),$$

so the ratio $\tilde{W}^h/W^h \equiv \tilde{\kappa}$ is common to all households; summing over households identifies it as

$$\tilde{\kappa} \;=\; \frac{M + q\tilde{B} + \theta(1-q)\tilde{B}}{M + qB + \theta(1-q)B}.$$

Let $(\bar{p}, q)$ with allocation $(\bar{x}^h)_h$ be a monetary equilibrium of $\mathcal{E}(\theta)$ at the stock B. Consider the stock $\tilde{B}$ at prices $(\tilde{\kappa}\bar{p}, q)$. For each household the collapsed constraint reads $\tilde{\kappa}\bar{p}\cdot\tau^{h,+} - q\,\tilde{\kappa}\bar{p}\cdot\tau^{h,-} \le \tilde{W}^h = \tilde{\kappa}\,W^h$; dividing by $\tilde{\kappa} > 0$ returns the constraint at $(\bar{p}, B)$. The set of attainable net trades is unchanged, so $\bar{x}^h$ remains optimal for every h, and goods markets clear. Financial positions supporting the allocation exist as in Steps 3 and 5 of the proof of Proposition 1, scaled to the new nominal volume, and the banking sector accommodates the implied aggregate credit at the unchanged policy rate. Hence $(\tilde{\kappa}\bar{p}, q)$ with allocation $(\bar{x}^h)_h$ extends to a monetary equilibrium of $\mathcal{E}(\theta)$ at the stock $\tilde{B}$. □

Proof of Proposition 3. Step 0: the two-unit constraint system. At prices $(p, q, \bar{\varepsilon})$, a plan for household h specifies net trades, a net financial position ψ^h as in the collapsed representation, a within-date conversion z^h (sovereign units acquired against coin at rate $\bar{\varepsilon}$; $z^h < 0$ is a sale), and a split of outlays across units: sovereign outlays $a^h \ge 0$ and coin outlays $a^h_S \ge 0$ with $a^h + a^h_S/\bar{\varepsilon} = p\cdot\tau^{h,+}$. The unit-level payment constraints are

$$a^h \;\le\; m^h - q\psi^h + z^h, \qquad a^h_S \;\le\; m^h_S - \bar{\varepsilon} z^h,$$

and at settlement all resources—sales receipts in either unit, unspent balances of either unit, and bond deliveries—discharge the obligation ψ^h, with coin counted at $1/\bar{\varepsilon}$:

$$\big[\text{sovereign receipts and balances}\big] + \tfrac{1}{\bar{\varepsilon}}\big[\text{coin receipts and balances}\big] + b^h + \psi^h \;\ge\; 0.$$

Conversion clears across households, $\sum_h z^h = 0$.

Step 1: aggregation. Add the two payment constraints, the coin constraint multiplied by $1/\bar{\varepsilon}$: the conversion terms cancel, and the sum reads $p\cdot\tau^{h,+} \le \tilde{m}^h - q\psi^h$ with $\tilde{m}^h \equiv m^h + m^h_S/\bar{\varepsilon}$— the baseline constraint (T) at the endowment $\tilde{m}^h$. The settlement constraint is likewise the baseline (R) at $\tilde{m}^h$. Conversely, any baseline plan at $(\tilde{m}^h, b^h)$ is implemented in the two-unit system with $z^h = 0$ by splitting every outlay and discharge across the units in proportion to holdings. The projection of the two-unit budget set onto consumption bundles therefore equals the baseline projection at $(\tilde{m}^h, b^h)$, and Lemma 1 and the collapsed representation apply at these endowments.

Step 2: existence for every $\bar{\varepsilon}$. The aggregated economy is the baseline of Section 2 with endowments $(\tilde{m}^h, b^h)$. The gains-from-trade measure $\gamma(e)$ depends only on real data, so $\gamma(e) > r$ delivers a monetary equilibrium $(\bar{p}, q)$ of the aggregated economy by the existence theorem (Remark 1). By Step 1, its allocation and prices, implemented pro rata with $z^h = 0$ and coin prices $\bar{p}^S = \bar{\varepsilon}\bar{p}$, constitute a monetary equilibrium with two outside monies: unit-level feasibility holds by construction, goods markets clear, aggregate credit is accommodated as in the baseline, and the unit-level flows balance because every payment and discharge divides between the units in proportion to holdings, so the two stocks circulate and settle pro rata, absorbing $(1-q)\bar{D} = \tilde{M} + B$ in sovereign-equivalent terms. This proves (i).

Step 3: scaling and indeterminacy. Under the proportionality hypothesis, $m_S^h = \mu\,(m^h + qb^h)$ with $\mu \equiv M_S/(M + qB)$, so effective liquid wealth is

$$\tilde{m}^h + qb^h \;=\; \big(1 + \mu/\bar{\varepsilon}\big)(m^h + qb^h),$$

a common factor across households for every $\bar{\varepsilon}$. By the scaling argument in the proof of Corollary 2(ii), the equilibria at different $\bar{\varepsilon}$ can therefore be constructed from a single one, with the same allocation and all sovereign prices scaled by the ratio of $M + M_S/\bar{\varepsilon} + qB$; coin prices follow by no-arbitrage. This proves (iii). For (ii): since $M_S > 0$, distinct $\bar{\varepsilon}$ yield distinct nominal price systems, so the construction produces a nondegenerate continuum of monetary equilibria indexed by $\bar{\varepsilon}$; and since $\bar{\varepsilon}$ enters the equilibrium conditions only through the aggregate $\tilde{M} = M + M_S/\bar{\varepsilon}$, no equilibrium condition excludes any $\bar{\varepsilon} > 0$. □

References

Adrian, T., P. Bains, M. Bechara, E. Cerutti, S. Forte, F. Grinberg, A. Gullo, M. Hengge, A. Jekabsone, K. Kao, T. M. Griffoli, S. M. Peria, M. Miccoli, M. Reuter, and N. Sugimoto (2025). *Understanding Stablecoins*. Tech. rep. URL: https://www.imf.org/en/-/media/files/publications/dp/2025/english/usea.pdf.

Adrian, T. and T. Mancini Griffoli (2019). *The Rise of Digital Money*. FinTech Notes. Washington, D.C: International Monetary Fund. ISBN: 978-1-4983-2490-8. DOI: 10.5089/9781498324908.063.

Ahmed, R. and I. Aldasoro (May 2025). *Stablecoins and Safe Asset Prices*. Working Paper 1270. Basel: Bank for International Settlements. URL: https://www.bis.org/publ/work1270.pdf (visited on 12/02/2026).

Anadu, K., P. Azar, M. Cipriani, T. M. Eisenbach, C. Huang, M. Landoni, G. La Spada, M. Macchiavelli, A. Malfroy-Camine, and J. C. Wang (2024). "Runs and Flights to Safety: Are Stablecoins the New Money Market Funds?" In: *Staff Reports* 1073. URL: https://www.newyorkfed.org/research/staff_reports/sr1073.html.

Andolfatto, D. (Mar. 2021). "Assessing the Impact of Central Bank Digital Currency on Private Banks". In: *The Economic Journal* 131.634, pp. 525–540. ISSN: 0013-0133, 1468-0297. DOI: 10.1093/ej/ueaa073. URL: https://academic.oup.com/ej/article/131/634/525/5900973 (visited on 08/19/2026).

Angeloni, I. (2026). *Whither Stablecoins?* Working Paper. URL: https://ignazioangeloni.home.blog/2026/05/05/working-paper-whither-stablecoins/ (visited on 07/16/2026).

Aronoff, D. J., F. C. Calabia, A. Brownworth, A. Samuel, and N. Narula (Feb. 2026). *The Hidden Plumbing of Stablecoins: Financial and Technological Risks in the GENIUS Act Era*. Tech. rep. MIT Digital Currency Initiative. URL: https://www.media.mit.edu/publications/the-hidden-plumbing-of-stablecoins-financial-and-technological-risks-in-the-genius-act-era/.

Azzimonti, M. and V. Quadrini (July 2025). *Digital Economy, Stablecoins, and the Global Financial System*. Tech. rep. w34066. Cambridge, MA: National Bureau of Economic Research, w34066. DOI: 10.3386/w34066. URL: http://www.nber.org/papers/w34066.pdf (visited on 01/16/2026).

Barro, R. J. (Nov. 1974). "Are Government Bonds Net Wealth?" In: *Journal of Political Economy* 82.6, pp. 1095–1117. ISSN: 0022-3808, 1537-534X. DOI: 10.1086/260266. URL: https://www.journals.uchicago.edu/doi/10.1086/260266 (visited on 08/22/2026).

Barthelemy, J., P. Gardin, and B. Nguyen (Jan. 2024). *Stablecoins and Short-Term Funding Markets*. SSRN Scholarly Paper. Rochester, NY. DOI: 10.2139/ssrn.4706063.

Social Science Research Network: 4706063. URL: https://papers.ssrn.com/abstract=4706063 (visited on 08/11/2026).

Barthelemy, J., J. Pedrono, and J. Schmidt (Dec. 2025). *The International Monetary System: Are Stablecoins Game Changers?* SSRN Scholarly Paper. Rochester, NY. DOI: 10.2139/ssrn.5953957. Social Science Research Network: 5953957. URL: https://papers.ssrn.com/abstract=5953957 (visited on 08/11/2026).

Benigno, P. (2026). *Stablecoins and Central Bank Digital Currencies: Who Supplies Liquidity?* Working Paper.

Bessent, S. (July 2025). *Statement on Enactment of the GENIUS Act*. URL: https://home.treasury.gov/news/press-releases/sb0197 (visited on 02/16/2026).

Bindseil, D. U. and E. Voloder (2026). *Reforming MiCA for Euro Stablecoins*. Tech. rep. URL: https://www.blockchain4europe.eu/wp-content/uploads/2026/04/Reforming-MiCA-for-Euro-Stablecoins-26426.pdf.

BIS (June 2025). *III. The next-Generation Monetary and Financial System*. BIS Annual Report. Basel: Bank for International Settlements, pp. 79–113. URL: https://www.bis.org/publ/arpdf/ar2025e3.htm (visited on 07/16/2026).

Boar, C. and A. Wehrli (2021). *Ready, Steady, Go? Results of the Third BIS Survey on Central Bank Digital Currency*. BIS Papers 114. Basel: Bank for International Settlements, Monetary and Economic Department.

Borio, C., P. Disyatat, and N. Tarashev (Feb. 2026). *Stablecoins and the Exorbitant Privilege of Money Creation*. URL: https://cepr.org/voxeu/columns/stablecoins-and-exorbitant-privilege-money-creation.

Born, A., C. Lambert, B. Nguyen, O. Soons, D. Staunton, and A. van der Kraaij (Apr. 2026). *Euro Stablecoins and Their Potential Effect on Sovereign Bond Markets*. Tech. rep. 33. European Central Bank. URL: https://www.ecb.europa.eu/press/financial-stability-publications/macroprudential-bulletin/html/ecb.mpbu202604_05.en.html (visited on 08/07/2026).

Brunnermeier, M. K. and D. Niepelt (Oct. 2019). "On the Equivalence of Private and Public Money". In: *Journal of Monetary Economics*. SPECIAL CONFERENCE ISSUE: "Money Creation and Currency Competition" October 19-20, 2018 Sponsored by the Study Center Gerzensee and Swiss National Bank 106, pp. 27–41. ISSN: 0304-3932. DOI: 10.1016/j.jmoneco.2019.07.004. URL: https://www.sciencedirect.com/science/article/pii/S0304393219301229 (visited on 08/19/2026).

Caixabank (2025). *CaixaBank and 9 European Banks to Launch Qivalis Euro Stablecoin*. URL: https://www.caixabank.com/en/headlines/news/qivalis-joint-venture-of-a-european-banking-consortium-to-launch-euro-stablecoin-in-the-second-half-of-2026 (visited on 08/07/2026).

Cecchetti, S. G. and K. L. Schoenholtz (June 2019). *Italeave: The Renewed Threat of a Parallel Currency*. URL: https://www.moneyandbanking.com/commentary/

2019/6/9/italeave-the-renewed-threat-of-a-parallel-currency (visited on 08/07/2026).

Cipollone, P. (Feb. 2026). “The Digital Euro - Strengthening Europe’s Payments Ecosystem”. URL: https://www.bis.org/review/r260217d.htm (visited on 08/07/2026).

Circle Internet Group (Apr. 2025). *Registration Statement on Form S-1*. Tech. rep. U.S. Securities and Exchange Commission. URL: https://www.sec.gov/Archives/edgar/data/1876042/000119312525070481/d737521ds1.htm (visited on 02/12/2026).

Circle Internet Group, Inc. (Oct. 2025). *USDC Reserve Report, September 2025*. Reserve Attestation Report. New York: Circle Internet Group, Inc. URL: https://6778953.fs1.hubspotusercontent-na1.net/hubfs/6778953/USDCAttestationReports/2025/2025%20USDC_Examination%20Report%20September%2025.pdf (visited on 02/12/2026).

Clower, R. W. (1967). “A Reconsideration of the Microfoundations of Monetary Theory”. In: *Western Economic Journal* 6.1, pp. 1–8.

Draghi, M. (June 2019). *Introductory Statement to the Press Conference (with Q&A)*. URL: https://www.ecb.europa.eu/press/press_conference/monetary-policy-statement/2019/html/ecb.is190606~32b6221806.en.html (visited on 08/07/2026).

Dubey, P. and J. Geanakoplos (Mar. 1992). “The Value of Money in a Finite-Horizon Economy: A Role for Banks”. In: *Economic Analysis of Markets and Games*. Ed. by P. Dasgupta, D. Gale, O. Hart, and E. Maskin. The MIT Press, pp. 407–444. ISBN: 978-0-262-27120-2. DOI: 10.7551/mitpress/2581.003.0022. URL: https://direct.mit.edu/books/book/4381/chapter/187565/The-Value-of-Money-in-a-Finite-Horizon-Economy-A (visited on 01/16/2026).

— (Mar. 2003). “Inside and Outside Fiat Money, Gains to Trade, and IS-LM”. In: *Economic Theory* 21.2-3, pp. 347–397. ISSN: 0938-2259, 1432-0479. DOI: 10.1007/s00199-002-0296-5. URL: http://link.springer.com/10.1007/s00199-002-0296-5 (visited on 01/16/2026).

— (Jan. 2006). “Determinacy with Nominal Assets and Outside Money”. In: *Economic Theory* 27.1, pp. 79–106. ISSN: 0938-2259, 1432-0479. DOI: 10.1007/s00199-004-0591-4. URL: http://link.springer.com/10.1007/s00199-004-0591-4 (visited on 01/16/2026).

— (Nov. 2010). “Credit Cards and Inflation”. In: *Games and Economic Behavior* 70.2, pp. 325–353. ISSN: 08998256. DOI: 10.1016/j.geb.2010.02.004. URL: https://linkinghub.elsevier.com/retrieve/pii/S0899825610000321 (visited on 01/16/2026).

Dubey, P. and J. Geanakoplos (2024). “Monies and Inflation in a Finite Horizon Economy”. In: *Klein Lecture*. University of Pennsylvania. URL: https://economics.sas.upenn.edu/system/files/2024-09/2024Paper.pdf.

— (Feb. 2026). “2024 Klein Lecture - Monies, Transactions, and Inflation: Part 1”. In: *International Economic Review* 67.1, pp. 3–33. ISSN: 0020-6598, 1468-2354. DOI: 10.1111/iere.70049. URL: https://onlinelibrary.wiley.com/doi/10.1111/iere.70049 (visited on 07/16/2026).

Eichengreen, B. (June 2025). “The Genius Act Will Bring Economic Chaos”. In: *The New York Times*. ISSN: 0362-4331. URL: https://www.nytimes.com/2025/06/17/opinion/genius-act-stablecoin-crypto.html (visited on 08/07/2026).

European Parliament and Council of the European Union (June 2023). *Regulation (EU) 2023/1114 on Markets in Crypto-Assets (MiCAR)*. URL: https://eur-lex.europa.eu/eli/reg/2023/1114/oj.

Fama, E. F. (1980). “Banking in the Theory of Finance”. In: *Journal of Monetary Economics* 6.1, pp. 39–57.

Fernández-Villaverde, J., D. Sanches, L. Schilling, and H. Uhlig (July 2021). “Central Bank Digital Currency: Central Banking for All?” In: *Review of Economic Dynamics* 41, pp. 225–242. ISSN: 10942025. DOI: 10.1016/j.red.2020.12.004. URL: https://linkinghub.elsevier.com/retrieve/pii/S1094202520301150 (visited on 08/19/2026).

Ferrari Minesso, M. and D. Siena (2026). *Private Money and Public Debt*. Tech. rep. 3174. DE: European Central Bank. DOI: 10.2866/1321559. URL: https://data.europa.eu/doi/10.2866/1321559 (visited on 07/16/2026).

Foley, D. K. (Sept. 1970). “Economic Equilibrium with Costly Marketing”. In: *Journal of Economic Theory* 2.3, pp. 276–291. ISSN: 00220531. DOI: 10.1016/0022-0531(70)90042-6. URL: https://linkinghub.elsevier.com/retrieve/pii/0022053170900426 (visited on 07/24/2026).

Gersbach, H., H. van Buggenum, and S. Zelzner (2026). *Stablecoins as Private Money: A Policy Agenda*. Tech. rep. 789. Center for Financial Studies.

Gersbach, H. and S. Zelzner (2025). *Stablecoin Competition in Rollover Markets: Yield Design and the Funds-Concentration Effect*. Tech. rep.

Gorton, G. and G. Pennacchi (1990). “Financial Intermediaries and Liquidity Creation”. In: *The Journal of Finance* 45.1, pp. 49–71. ISSN: 1540-6261. DOI: 10.1111/j.1540-6261.1990.tb05080.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.1990.tb05080.x (visited on 08/19/2026).

Gorton, G. and J. Zhang (2023). “Taming Wildcat Stablecoins”. In: *The University of Chicago Law Review* 90.3. URL: https://chicagounbound.uchicago.edu/uclrev/vol90/iss3/3/.

Gurley, J. G. and E. S. Shaw (1960). *Money in a Theory of Finance*. Washington, DC: Brookings Institution.

Hayek, F. A. (1976). *Denationalisation of Money*. Hobart Paper 70. London: Institute of Economic Affairs.

Hellwig, M. (1985). "What Do We Know about Currency Competition?" In: *Zeitschrift für Wirtschafts- und Sozialwissenschaften* 105.5, pp. 565–588.

Hofmann, B., M. Kaldorf, and M. Rottner (2026a). *The Macroeconomics of Stablecoins*. BIS Working Paper 1363. Bank for International Settlements. URL: https://www.bis.org/publ/work1363.htm.

Hofmann, B., A. Mehrotra, and J. Paulick (2026b). *Dollarisation and Monetary Control: What Lessons for the Rise of Stablecoins?* BIS Working Paper 1370. Bank for International Settlements.

Kareken, J. and N. Wallace (May 1981). "On the Indeterminacy of Equilibrium Exchange Rates". In: *The Quarterly Journal of Economics* 96.2, pp. 207–222. ISSN: 00335533. DOI: 10.2307/1882388. URL: https://academic.oup.com/qje/article-lookup/doi/10.2307/1882388 (visited on 07/29/2026).

Keister, T. and D. Sanches (Jan. 2023). "Should Central Banks Issue Digital Currency?" In: *The Review of Economic Studies* 90.1, pp. 404–431. ISSN: 0034-6527, 1467-937X. DOI: 10.1093/restud/rdac017. URL: https://academic.oup.com/restud/article/90/1/404/6561552 (visited on 08/19/2026).

Krishnamurthy, A. and A. Vissing-Jorgensen (Apr. 2012). "The Aggregate Demand for Treasury Debt". In: *Journal of Political Economy* 120.2, pp. 233–267. ISSN: 0022-3808, 1537-534X. DOI: 10.1086/666526. URL: https://www.journals.uchicago.edu/doi/10.1086/666526 (visited on 08/19/2026).

Lagarde, C. (May 2026). *Stablecoins and the Future of Money: Separating Functions from Instruments*. URL: https://www.ecb.europa.eu/press/key/date/2026/html/ecb.sp260508~dd909fbed1.en.html (visited on 08/07/2026).

Li, B., T. Mancini Griffoli, M. Miccoli, B. Tan, and L. Zhang (2026). *Making Stablecoins Stable*. Tech. rep. 26/74. IMF.

Libra Association (Dec. 2019). *An Introduction to Libra: White Paper*. URL: https://web.archive.org/web/20191228052721/https://libra.org/en-US/white-paper/ (visited on 08/11/2026).

Ma, Y., Y. Zeng, and A. L. Zhang (2025). "Stablecoin Runs and the Centralization of Arbitrage". In: URL: https://bfi.uchicago.edu/insights/stablecoin-runs-and-the-centralization-of-arbitrage/.

Magill, M. and M. Quinzii (1996). *Theory of Incomplete Markets*. Vol. 1. Cambridge, MA: MIT Press.

Mehrling, P. (2013). "The Inherent Hierarchy of Money". In: Routledge.

Mehrling, P. (2020). *Payment vs. Funding: The Law of Reflux for Today*. INET Working Paper 113. Institute for New Economic Thinking. URL: https://www.ineteconomics.org/uploads/papers/WP_113-Mehrling-Payment-vs-Funding.pdf.

Monetary authority of Singapore (Oct. 2022). *MAS Project Orchid Report*. Tech. rep. Monetary Authority of Singapore. URL: https://www.mas.gov.sg/-/media/mas-media-library/development/fintech/project-orchid/mas-project-orchid-report.pdf.

Niepelt, D. (2026). *Central Bank Digital Currency and Monteary Architecture*. Tech. rep. DP21141. CEPR.

Odinet, C. K. and A. Tosato (2025). "Regulating Stablecoins: Comparing MiCAR and the GENIUS Act". In.

Schaaf, J. (July 2025). *From Hype to Hazard: What Stablecoins Mean for Europe*. URL: https://www.ecb.europa.eu/press/blog/date/2025/html/ecb.blog20250728~e6cb3cf8b5.en.html (visited on 08/07/2026).

Shubik, M. (1999). *The Theory of Money and Financial Institutions*. Cambridge, MA: MIT Press.

Tether International, S.A. de C.V. (2025). *Financial Figures and Reserve Report, Q3 2025*. Independently Attested Reserves Report. El Salvador: Tether International, S.A. de C.V. URL: https://assets.ctfassets.net/vyse88cgwfbl/6GbUTVK4tTYAytefu5daIi/6cac18eb4b526c9c52640a3d2bed9642/ISAE_3000R_-_Opinion_Tether_International_Financial_Figure_31-10-2025.pdf (visited on 02/12/2026).

Tobin, J. (1963). "Commercial Banks as Creators of "Money"". In: *Banking and Monetary Studies*, pp. 408–419.

U.S. Congress (July 2025). *Guiding and Establishing National Innovation for U.S. Stablecoins Act of 2025 (GENIUS Act)*. URL: https://www.congress.gov/bill/119th-congress/senate-bill/1582 (visited on 02/16/2026).

UBS Switzerland (2026). *UBS, PostFinance, Sygnum, Raiffeisen, Zürcher Kantonalbank, BCV and Swiss Stablecoin AG Launch Joint CHF Stablecoin Sandbox | UBS Global*. URL: https://www.ubs.com/global/en/media/display-page-ndp/en-20260408-stablecoin.html (visited on 08/07/2026).

Van Buggenum, H., H. Gersbach, and S. Zelzner (2024). "Contagious Stablecoins?" In: *SSRN Electronic Journal*. ISSN: 1556-5068. DOI: 10.2139/ssrn.4746634. URL: https://www.ssrn.com/abstract=4746634 (visited on 04/08/2026).

Waller, C. J. (Aug. 2021). "CBDC - A Solution in Search of a Problem?" URL: https://www.bis.org/review/r210806a.htm (visited on 08/11/2026).

The Working Paper series of the Oesterreichische Nationalbank is designed to disseminate and to provide a platform for discussion of either work of the staff of the OeNB economists or outside contributors on topics which are of special interest to the OeNB. To ensure the high quality of their content, the contributions are subjected to an international refereeing process. The opinions are strictly those of the authors and do in no way commit the OeNB.

The Working Papers are also available on our website (http://www.oenb.at) and they are indexed in RePEc (http://repec.org/).

Publisher and editor Oesterreichische Nationalbank
Otto-Wagner-Platz 3, 1090 Vienna, Austria
PO Box 61, 1011 Vienna, Austria
www.oenb.at
oenb.info@oenb.at
Phone (+43-1) 40420-6666

Editor Martin Summer

Cover Design Information Management and Services Division

DVR 0031577

ISSN 2310-533X (Online)