When Stablecoins Meet Treasury Scarcity

Global Spillovers from Dollar Settlement

Lorenzo Forni [1,2,*]

[1]University of Padua, Department of Economics
[2]Prometeia Associazione

June 2, 2026

Abstract

We study how USD-denominated stablecoins affect international macroeconomic transmission when they become the dominant global settlement technology. We develop a two-country DSGE model in which transactions must be settled in USD stablecoins backed by short-term US Treasury securities. Because backing capacity is limited by the supply of safe dollar assets, increases in USD settlement demand raise the marginal cost of collateral and generate an endogenous transaction fee that is increasing and convex in safe dollar assets utilization. We show that this feature creates a new international spillover channel. Quantitatively, we calibrate three Treasury-scarcity regimes that produce annualised steady-state settlement fees of 50 (baseline regime), 150 (stressed) and 300 (extreme) basis points. Moving across regimes generates large *permanent* welfare effects: Foreign consumption falls by 0.35% and 0.87% in the stressed and extreme regimes respectively, with consumption-equivalent welfare losses of 0.51% and 1.27%. We then extend the model to include a foreign central bank digital currency (CBDC). Although dollar stablecoins remain dominant because of dollar invoicing and network effects, CBDC provides state-dependent insurance: at the extreme regime, it reduces the steady-state stablecoin fee by 62% and the Foreign welfare loss by 79%; around the baseline regime, where the fee schedule is comparatively flat, shocks generate fee deviations of at most ±1.2 basis points and CBDC has little effect; the insurance value of CBDC is strongly convex in the severity of Treasury scarcity. Our results imply that the relevant financial-stability metric in a stablecoin-based international payments system is not stablecoin market capitalization alone, but stablecoin backing demand relative to the supply of short-term US Treasuries.

JEL Codes: F31, F33, F41, G15, E42

Keywords: Stablecoins, Safe Assets, International Finance, CBDC, Dollar Dominance, Settlement Systems

---

[*] Department of Business and Economics, University of Padua, Via del Santo 33, Padova, Italy and Prometeia Associazione, Piazza Trento e Trieste 3, Bologna, Italy. E-mail: `lorenzo.forni@unipd.it`. This study is financed by the European Union – NextGenerationEU, in the project PRIN2022 (DM 104 del 02.02.2022) "SustainablE LEgaCy DebT (SELECT)", project code 2022KCEYML, CUP C53D23002790006.

1 Introduction

The US dollar occupies a unique position in the global financial system, serving simultaneously as the world's dominant medium of exchange, unit of account, and store of value. This "exorbitant privilege" manifests in multiple dimensions: the dollar accounts for approximately 60% of global foreign exchange reserves, 50% of international trade invoicing, and 90% of foreign exchange transactions (Gopinath and Stein, 2021; Maggiori et al., 2020). A central pillar supporting this dominance is the depth and liquidity of US Treasury markets, which provide the benchmark safe asset for global investors. The convenience yield on US Treasuries — reflecting their unparalleled safety and liquidity — allows the United States to borrow at lower interest rates than fundamentals alone would justify (Krishnamurthy and Vissing-Jorgensen, 2012; Jiang et al., 2021).

The emergence of blockchain-based stablecoins introduces a new dimension to this system. Stablecoins are cryptocurrencies designed to maintain a stable value relative to fiat currencies, most commonly pegged 1:1 to the US dollar. The largest stablecoins — Tether (USDT), USD Coin (USDC), and others — are backed primarily by short-term US Treasury securities. By end-2025, the combined market capitalisation of USD-denominated stablecoins exceeded $300 billion, with the two largest issuers holding over $200 billion in direct and indirect Treasury exposure — a position that places them collectively among the world's largest holders of US government debt, ahead of Germany and Saudi Arabia.

The GENIUS Act, signed into law in July 2025 (Public Law 119–27), established the first federal regulatory framework for US payment stablecoins and formalised the link between stablecoin issuance and safe dollar assets. The Act requires permitted issuers to maintain reserves on at least a 1:1 basis with outstanding tokens, drawn from an enumerated list of eligible assets. This list includes: US coins and currency and Federal Reserve account balances; demand deposits at insured depository institutions; Treasury bills, notes, or bonds with a remaining maturity of 93 days or less; overnight repurchase agreements backed by Treasury securities with maturities of 93 days or less; overnight reverse repurchase agreements collateralised by Treasury securities; and government money market funds invested solely in the foregoing. Any of these may be held in tokenised form. Reserve assets may not be rehypothecated except to create short-term liquidity for redemption purposes. While the eligible reserve universe is therefore broader than Treasury bills alone — encompassing cash, central bank deposits, and insured bank deposits — the economics of stablecoin issuance strongly favour interest-bearing assets. Under the Act, issuers are explicitly prohibited from paying any interest or yield to stablecoin holders. Reserve income thus accrues entirely to issuers as profit, creating a powerful incentive to maximise the yield on reserves. Cash and demand deposits at commercial banks typically earn interest rates that are lower than, or at most comparable to, the federal funds rate, while short-term Treasury bills offer market rates with no credit risk. Government money market funds, themselves primarily invested in Treasury securities and Treasury repo, similarly earn near-Treasury rates. As a result, profit-maximising issuers have a strong structural incentive to concentrate reserves in short-term Treasuries and Treasury-backed repo rather than in cash or deposits, even when the latter are technically permissible. This is borne out in practice: as of Q4 2025, Tether held approximately USD 122 billion in direct Treasury bills and USD 18.7 billion in Treasury repo out of USD 192.9 billion in total reserves, while Circle held approximately USD 24.5 billion in direct Treasuries and USD 40.6 billion in Treasury repo out of USD 76.5 billion — implying that over 85% of both issuers' reserves were effectively invested in Treasury instruments. This framework is likely to increase structural demand for Treasury bills, although the scale, persistence, and elasticity of that demand remain market-dependent.

This paper asks: *How does the adoption of Treasury-backed stablecoins as a global settlement technology affect international macroeconomic dynamics, and what role can central bank digital currencies (CBDCs) play in reshaping this system?*

We develop a quantitative two-country dynamic stochastic general equilibrium (DSGE) model where stablecoins are not merely portfolio assets but *essential settlement infrastructure* subject to a binding capacity constraint. In our model all transactions must be settled using USD-denominated stablecoins, whose aggregate supply is limited by the available stock of short-term US Treasuries (*backing capacity*). When settlement demand approaches backing capacity, an endogenous convex fee arises — capturing congestion in collateral markets akin to repo market stress during Treasury scarcity. The fee arises because stablecoin issuers must continuously source, fund, and roll short-term Treasury collateral to back their liabilities. When settlement demand is low relative to the available stock of Treasuries, this is nearly costless: issuers can access collateral at market rates with minimal friction. As demand grows relative to backing capacity, however, the marginal cost of accommodating additional settlement rises sharply. Issuers must compete for increasingly scarce bills, pay up in repo markets, manage larger intraday liquidity buffers, and roll collateral more frequently — frictions that are inherently non-linear and that issuers pass through to users as a transaction fee. We capture this with a fee that is an increasing and convex function of the utilisation rate — the ratio of settlement demand to backing capacity. At low utilisation the fee is nearly flat and settlement is almost costless. As utilisation approaches the backing limit, the fee rises steeply and without bound, rationing demand.

This generates a "safe asset capacity channel": shocks to US Treasury bill supply propagate internationally through settlement costs, creating spillovers distinct from traditional trade, valuation, or portfolio balance channels. In our model, the Home producer price is fixed at one — the dollar is the unit of account — but the Foreign producer price (in FCU) moves endogenously: it rises when Foreign wages increase relative to Foreign productivity, and falls when Foreign productivity improves or when Foreign output expands relative to employment. The Foreign good price expressed in USD, $P_{FH,t} = P_{FF,t}/E_t$, additionally falls when the Foreign currency depreciates against the dollar. Because Foreign stablecoin settlement demand depends on the dollar value of Foreign invoices — which depends on Foreign prices divided by the exchange rate — any factor that lowers Foreign prices in dollar terms (a Foreign productivity boom, a Foreign currency depreciation, or an expansion of Foreign output relative to employment) directly reduces settlement demand, independently of what is happening to trade volumes. This means that any such development acts as a partial release valve for the global settlement system. By contrast, a US productivity boom raises US output and dollar income without any offsetting compression of invoice values, so it pushes settlement demand upward more forcefully. This asymmetry — Foreign cost conditions can relieve settlement pressure while US conditions cannot — means that Foreign shocks are systematically less congestive for the global settlement system than US shocks of equal magnitude.

Main results. Our central finding concerns the *permanent* welfare implications of operating the global settlement system at different levels of Treasury scarcity. We calibrate three regimes that produce annualised steady-state stablecoin fees of 50 (baseline), 150 (stressed) and 300 (extreme) basis points. Moving from the baseline to the stressed and extreme regimes lowers Foreign consumption permanently by 0.35% and 0.87%, with consumption-equivalent welfare losses of 0.51% and 1.27%. The cross-border fee transfer from Foreign to Home — a *digital exorbitant privilege* that mirrors the classical version through fees on settlement infrastructure rather than yield differentials — rises from 0.5% of Foreign GDP annually at the baseline regime to 1.5% (stressed) and 2.9% (extreme).

Around each steady state, the safe asset capacity channel is a *threshold phenomenon.* At the baseline regime ($u_{ss} \approx 0.54$) the convex fee schedule is nearly flat: shocks generate fee deviations of at most ± 1.2 basis points and real effects are dominated by standard RBC channels. At the extreme regime ($u_{ss} \approx 0.82$) the slope of the fee schedule is roughly eightfold steeper, so the same shocks activate three reinforcing channels — wealth transfer, intertemporal distortion, and congestion feedback — producing first-order cross-country spillovers.

A foreign CBDC provides state-dependent insurance that is strongly convex in Treasury scarcity. It reduces the steady-state stablecoin fee by 38%, 51% and 62% across the three regimes, and the Foreign consumption-equivalent welfare loss by 71% (stressed) and 79% (extreme): a nonlinear circuit-breaker that is largely irrelevant at the baseline and becomes decisive under stress.

These results contribute to three strands of literature. First, we provide the first quantitative general-equilibrium analysis of stablecoin adoption with both permanent (cross-steady-state) and dynamic (impulse-response) margins, identifying the ratio of stablecoin backing demand to short-term Treasury supply as the key structural parameter. Second, we formalise the digital exorbitant privilege as a permanent income flow generated by ownership of the dominant settlement layer rather than by yield differentials. Third, we characterise the conditions under which a foreign CBDC delivers material macroeconomic stabilisation — conditions concentrated in stressed and extreme regimes of safe-asset scarcity.

Quantitative summary. The cross-steady-state landscape summarised in Table 3 is the headline quantitative result. Foreign consumption is permanently 0.35% (stressed) and 0.87% (extreme) lower than at the baseline; the cross-border fee transfer reaches 1.5% and 2.9% of Foreign GDP annually; and the Foreign consumption-equivalent welfare loss is 0.51% and 1.27%. CBDC reduces the welfare-equivalent loss by 71% (stressed) and 79% (extreme), reflecting the convex insurance value of an alternative settlement instrument that absorbs marginal demand from a congested layer. The IRF analysis at the baseline regime (Section 6.2) confirms that shocks generate small fee deviations of at most ± 1.2 basis points (annualised) and negligible real effects; deviations scale up roughly with the slope of the fee schedule when one moves to the stressed and extreme regimes, completing the threshold story summarised above.

1.1 Related literature

Our work contributes to three strands of literature.

Stablecoins and digital currencies. Recent theoretical work has analysed stablecoins through the lens of banking and monetary theory. Gorton and Zhang (2023) compare stablecoins to private bank liabilities, emphasising the role of backing assets in ensuring convertibility. Catalini and de Gortari (2021) highlight stablecoins' potential to reduce cross-border payment frictions in emerging markets. Cong et al. (2022) and Routledge and Zetlin-Jones (2022) examine fragility and run risk in algorithmic stablecoins. Empirical studies by Lyons and Viswanath-Natraj (2023) and Makarov and Schoar (2020) document arbitrage dynamics and the role of stablecoins as vehicles for moving capital across cryptocurrency exchanges.

Most closely related to our work, Ferrari Minesso and Siena (2026) provide empirical evidence that major stablecoin issuers hold approximately $200 billion in US Treasuries and develop a portfolio model showing how stablecoin adoption tightens the global safe asset market. Azzimonti and Quadrini (2025) develop a three-region model — the United States, the rest of the world, and a non-geographic "Digital Economy" — in which stablecoins are explicitly treated as a store of value rather than as a means of payment. The diffusion of the Digital Economy activates two opposing channels: a "financial demand" channel that raises global demand for safe dollar assets (through stablecoin reserve backing) and a "real demand" channel that shifts production toward digitally-supplied services. Based on their model, the financial demand channel dominates in the long run, producing lower US interest rates, higher US foreign borrowing, and riskier portfolios both in the US and abroad. Ahmed and Aldasoro (2025) estimate that stablecoin flows have a statistically significant effect on Treasury bill yields, with magnitudes comparable to central bank reserve flows.

Our contribution is to move beyond portfolio allocation to model stablecoins as *transaction-essential settlement infrastructure* with endogenous, state-dependent fees arising from capacity constraints. This generates the threshold transmission mechanism — and the asymmetric revaluation and cost-deflation channels — that are absent from portfolio-based frameworks.

CBDCs and currency competition. The CBDC literature explores implications for bank intermediation (Fernández-Villaverde et al., 2021; Brunnermeier and Niepelt, 2019), monetary policy transmission (Davoodalhosseini, 2022), and international currency competition (Benigno et al., 2022; Auer et al., 2021). Barrdear and Kumhof (2022) study a retail CBDC issued against government debt in a closed-economy DSGE, finding that CBDC stocks on the order of 30% of GDP raise output substantially through lower real rates, reduced distortionary taxes, and lower transaction costs. Cova et al. (2026) develop a closely related two-country New Keynesian framework with cash, a globally-used stablecoin and a home CBDC, and show that monetary policy transmission depends on the composition of stablecoin backing: transmission is largely preserved when stablecoins are fully cash-backed but weakened when backing combines cash and government bonds, because the stablecoin-issuing fund rebalances its portfolio in response to interest rate changes. Their result parallels ours in emphasising backing composition as the key determinant of macroeconomic outcomes, but our mechanism is structurally distinct: we work with flexible prices and study settlement-cost spillovers rather than nominal monetary transmission, and we model the backing constraint as an endogenous, convex fee on stablecoin issuance rather than as a portfolio-rebalancing wedge. Benigno et al. (2022) show that when a foreign CBDC offers greater price stability and lower transaction costs than the domestic currency, households and firms may voluntarily adopt it as their preferred means of payment — a digital form of currency substitution analogous to traditional dollarisation but with much lower switching costs. As adoption spreads through network effects, the domestic central bank faces a shrinking monetary base and weakened policy transmission, creating an "impossible trinity": a country cannot simultaneously maintain exchange rate stability, independent monetary policy, and an open digital currency market. Avoiding monetary sovereignty loss requires either restricting access to foreign digital currencies or abandoning exchange rate and monetary policy targets. Le et al. (2023) analyse digital dollarisation in small open economies when foreign stablecoins are adopted.

Our CBDC extension differs by focusing on *coexistence* rather than displacement, and by characterising the insurance value of CBDC as a nonlinear function of backing capacity. We model competition between private, capacity-constrained stablecoins and public, elastically supplied CBDC, showing that invoicing conventions and network effects sustain stablecoin dominance under normal conditions — but that CBDC becomes a quantitatively important shock absorber when those conditions deteriorate. This speaks directly to the policy question of whether CBDC issuance can reduce dollar dominance: our answer is negative in normal times but qualified in stressed times.

International finance and safe assets. Our safe asset capacity channel builds on the literature emphasising global safe asset scarcity (Caballero et al., 2008; Caballero and Farhi, 2018), the dollar's exorbitant privilege (Gourinchas and Rey, 2007; Gourinchas et al., 2019), and the role of US Treasuries in international liquidity provision (He et al., 2019; Krishnamurthy and Vissing-Jorgensen, 2012). Maggiori et al. (2020) document dominant currency pricing in trade invoicing, showing the dollar's outsized role creates spillovers from US conditions. Gopinath and Stein (2021) formalise the international price system under dominant currency pricing.

We contribute by showing how tokenisation transforms Treasuries from portfolio stores of value into *transaction infrastructure* with a nonlinear congestion technology. When stablecoins back their liabilities with Treasuries and those stablecoins become essential for settlement, Treasury scarcity no longer operates solely through convenience yields and portfolio rebalancing — it directly affects the marginal cost of executing transactions worldwide, with effects that are dormant at low utilisation but become first-order as utilisation crosses the congestion threshold.

The cross-border fee transfer that emerges from this structure constitutes what we call a *digital exorbitant privilege*. Whereas the classical exorbitant privilege of Gourinchas and Rey (2007); Gourinchas et al. (2019) arises from a yield differential on the US external balance sheet — a higher return on foreign assets than is paid on foreign liabilities — the privilege identified in our model arises from ownership of the dominant settlement infrastructure. Three reinforcing channels operate within this digital privilege. First, Foreign agents pay a perpetual cross-border fee to use the USD-denominated payment rails owned by US-based stablecoin issuers, generating a one-way income flow to Home that does not require any return differential or net asset position. Second, the Treasury collateral required to back USD stablecoins generates additional demand for short-dated US sovereign debt, depressing yields and lowering Home's cost of borrowing

— a digital counterpart to the low-rate-on-liabilities side of the classical privilege. Third, the transactional demand for dollars that USD stablecoins create translates into FX-market demand for USD and tends to appreciate the dollar, generating a terms-of-trade transfer to Home. Our model formalises only the first of these channels; the second and third operate in the same direction and would reinforce our quantitative findings if endogenised. The two privileges — classical and digital — are structurally analogous and complementary rather than substitutes, and the digital channel may sustain Home's external income even in steady states where the classical channel has weakened.

1.2 Roadmap

The paper proceeds as follows. Section 2 documents the growth of USD stablecoins, their Treasury backing, and the mechanics of collateral-backed issuance. Section 3 presents the baseline model with stablecoins as essential settlement infrastructure. Section 4 extends the model to include foreign CBDC, deriving coexistence conditions and competitive discipline mechanisms. Section 5 discusses parameter values. Section 6 presents the quantitative results: first the cross-steady-state welfare comparison across the three Treasury-scarcity regimes (Section 6.1), then the impulse-response analysis around the baseline regime (Section 6.2), followed by a discussion of the CBDC issuance mechanism (Section 6.3). Section 7 reports the sensitivity of the model to the remaining structural parameters. Section 8 concludes. Computational details and detailed shock-by-shock IRF commentary are collected in Appendices A and B.

2 Stylised facts and motivation

This section documents three empirical regularities that motivate our modelling framework: (i) the rapid growth and scale of USD-denominated stablecoins, (ii) their structural reliance on US Treasury backing institutionalised by the GENIUS Act, and (iii) the mechanics of collateralised issuance that generate capacity constraints and their implications for international transmission.

2.1 The rise of USD stablecoins

Stablecoins are cryptocurrencies designed to maintain a stable value relative to a reference asset, most commonly the US dollar. Unlike unbacked cryptocurrencies such as Bitcoin or Ethereum, whose prices fluctuate freely based on supply and demand, stablecoins use various mechanisms — reserve backing, algorithmic stabilisation, or over-collateralisation — to peg their value at $1.00 per token.

The market capitalization of the largest USD-pegged stablecoins — Tether (USDT), USD Coin (USDC), Sky Dollar (USDS, formerly DAI), and Ethena's USDe — reveals several patterns from 2020 through 2025.[1]

Rapid growth. Total stablecoin market capitalisation grew from approximately $10 billion in early 2020 to over $300 billion by 2025. This growth vastly exceeds that of traditional payment systems and implies that stablecoin issuers now hold Treasury positions comparable to those of major sovereign investors.

---

[1] Data from CoinGecko and authors' calculations.

Dominance of fiat-backed designs. Tether and USD Coin — both backed primarily by short-term safe assets — account for over 90% of stablecoin market capitalisation. Algorithmic stablecoins such as TerraUSD (which collapsed in May 2022) and collateralized designs like Sky Dollar remain niche. This revealed preference for fiat backing reflects user demand for transparency and redemption guarantees, and it is this design that creates the structural link between stablecoin growth and Treasury demand that motivates our model.

Peg stability. The market price of major stablecoins has remained within ±1% of the $1.00 peg in recent years despite significant volatility in underlying cryptocurrency markets. During the March 2020 COVID crisis and the May 2022 Terra collapse, Tether and USDC were able to (almost) maintain their pegs, suggesting that reserve backing is effective at ensuring stability. An important exception occurred in March 2023, when USDC dropped to approximately $0.88 after Circle could not immediately withdraw reserves held at Silicon Valley Bank — an episode illustrating that reserve-quality constraints can themselves become a source of fragility.

Use cases beyond crypto. While stablecoins originated as trading vehicles within cryptocurrency exchanges, their use is expanding to include remittances, cross-border payments, and decentralised finance (DeFi) applications. On-chain data indicate that a growing share of stablecoin transactions involve real economic activity — goods and services payments, payroll, and trade finance — rather than purely speculative trading. It is this extension into genuine settlement functions that motivates modelling stablecoins in this paper as transaction-essential infrastructure rather than optional portfolio assets.

2.2 Treasury backing and the safe asset link

A defining feature of the largest stablecoins is their structural reliance on US Treasury securities. Table 1 summarises the reserve composition of Tether and Circle (issuer of USDC) as of Q4 2025, based on their public attestation reports.

Treasury dominance. Including both direct holdings and Treasury-collateralised reverse repurchase agreements, Tether's total US government debt exposure reached approximately $141 billion at end-2025, while Circle held an estimated $65 billion in Treasuries and Treasury repo.[2] Combined, the two issuers hold over $200 billion in direct and indirect US Treasury exposure, placing them collectively among the world's largest holders of US government debt. Industry forecasts suggest stablecoin issuers could collectively hold $400–500 billion in Treasuries by 2028, which would place them among the top five foreign holders of US debt and potentially rival traditional sovereign buyers such as China and Japan, both of which have been reducing their Treasury exposure.

Notably, the two issuers exhibit contrasting reserve strategies that are informative about the backing constraint modelled in this paper. Circle maintains a pure safe-asset reserve composition — 100% of USDC reserves are held in cash, US Treasuries, and overnight Treasury reverse repurchase agreements. Tether has diversified approximately 25% of its reserves into gold,

---

[2] Tether's $141bn figure is reported in its Q4 2025 attestation. For Circle, the Federal Reserve Board reports that as of August 2025, approximately 34% of USDC reserves were held in direct US Treasuries and 51% in Treasury reverse repurchase agreements; applying these shares to end-2025 reserves of $76.5bn yields approximately $65bn in combined Treasury exposure. See Federal Reserve FEDS Notes, "Banks in the Age of Stablecoins," December 2025.

Table 1: Reserve composition of major stablecoin issuers (Q4 2025)

Asset Class	**Tether ($bn)**	**Circle ($bn)**
US Treasury bills (direct)	122.3	24.5
Reverse repo agreements	18.7	40.6
Money market funds	—	—
Cash and bank deposits	3.5	10.8
Precious metals (gold)	17.5	0.0
Bitcoin	8.4	0.0
Secured loans and other	22.5	0.0
Total reserves	192.9	76.5
Tokens outstanding	186.5	75.3
Reserve ratio	103%	102%

Sources: Tether Financial Figures and Reserves Report (December 31, 2025), attested by BDO; Circle Q4 2025 Earnings Report and USDC Reserve Attestation (December 2025), attested by Deloitte & Touche. For Tether, direct US Treasury holdings of $122.3bn plus overnight and term reverse repurchase agreements collateralised by Treasuries bring total direct and indirect Treasury exposure to approximately $141bn. Tether's "Secured loans and other" category includes $17.0bn in over-collateralised secured loans, corporate bonds, and other investments. Tether's proprietary investment portfolio (AI, energy, fintech, etc.), valued at over $20bn, is fully segregated from USDT reserves and is *not* included in the figures above. For Circle, the majority of reserves are held in the Circle Reserve Fund (USDXX), an SEC-registered 2a-7 government money market fund managed by BlackRock, holding short-dated US Treasuries and overnight Treasury reverse repurchase agreements. Reserves exceed outstanding tokens due to accumulated interest income and, in Tether's case, unrealised gains on gold and Bitcoin holdings.

Bitcoin, and secured loans, assets that are costlier to fund and less liquid than Treasuries. Our model's capacity constraint reflects Circle's design more closely: when all reserves must be held in short-term safe USD assets, stablecoin issuance capacity is directly bounded by the available stock of such assets.

Why Treasuries? The preference for Treasury backing reflects three factors. First, *regulatory mandate*: the GENIUS Act, signed into law in July 2025, established the first federal regulatory framework for US stablecoins, requiring all payment stablecoin issuers to maintain 100% reserve backing in cash, short-dated Treasury securities, overnight repurchase agreements, or other approved government assets. This codified what had previously been industry practice into a legal requirement, making the link between stablecoin growth and Treasury demand a permanent structural feature of the market. Second, *liquidity*: Treasuries are the deepest and most liquid market globally, allowing issuers to scale rapidly without market impact — a critical feature given that Tether alone issued nearly $50 billion in new USDT during 2025. Third, *yield*: short-term Treasuries earn interest income that accrues entirely to issuers, as the GENIUS Act prohibits stablecoin issuers from paying interest to holders. This revenue stream is substantial: Tether reported over $10 billion in net profit for 2025, while Circle reported $770 million in Q4 2025 revenue alone, the vast majority from reserve income.

Structural demand and the capacity constraint. Unlike traditional Treasury holders who optimise portfolio allocation based on risk-return trade-offs, stablecoin issuers hold Treasuries for *technological reasons*: backing liabilities requires safe, liquid collateral. This creates inelastic demand — when stablecoin adoption grows, Treasury holdings must grow proportionally. Ahmed and Aldasoro (2025) estimate that stablecoin flows have a statistically significant effect on Treasury bill yields, with magnitudes comparable to central bank reserve flows. Ferrari Minesso and Siena (2026) estimate that each $1 billion increase in stablecoin market cap translates to \$800–900 million in Treasury demand, a pass-through far higher than for traditional financial intermediaries.

The key implication for our model is that this inelastic demand is not merely a portfolio phenomenon: it creates a *capacity constraint* on the global settlement system. When stablecoin adoption grows without a commensurate increase in short-term Treasury supply, the ratio of settlement demand to backing capacity rises. Once this ratio crosses approximately 70–75% utilisation, the convex fee schedule transitions from a flat regime — where settlement is nearly costless and macroeconomic spillovers are negligible — to a congested regime where small increases in demand generate large fee spikes and first-order international spillovers (see Section 3.4 below). Monitoring this ratio is therefore a natural financial stability indicator for the stablecoin era, complementary to traditional metrics of bank leverage or money market fund fragility.

2.3 Implications for international transmission

The stylised facts documented above suggest three channels through which stablecoin adoption reshapes international macroeconomic linkages. The first operates continuously; the second and third are threshold-dependent.

Direct safe asset channel. Stablecoin-driven demand for US Treasuries lowers US borrowing costs (convenience yield effect) and tightens the global safe asset market. This traditional portfolio channel has been analysed by Caballero et al. (2008) and Gourinchas et al. (2019) in the context of foreign official reserve accumulation; stablecoins extend this mechanism to private digital finance. It bears emphasis that this channel operates through the revealed preference of profit-maximising stablecoin issuers rather than through strict regulatory requirements. The GENIUS Act permits issuers to hold reserves in several non-Treasury forms — including demand deposits and Federal Reserve balances — but these instruments typically carry lower yields. As long as Treasury bills offer a meaningful spread over insured deposits, issuers have an incentive to tilt their portfolios toward Treasuries. This spread-driven demand is more elastic than the inelastic institutional Treasury demand assumed in some portfolio models: it would attenuate if Treasury yields compressed toward deposit rates, introducing a margin of self-correction that our model, which treats Treasury backing as the unique and exclusive collateral form, does not capture.

Settlement cost channel (novel, threshold-dependent). When stablecoins become essential for settling transactions — as they increasingly are in cryptocurrency markets, cross-border payments, and, in the extreme scenario modelled here, domestic transactions as well — capacity constraints can generate endogenous transaction costs. Shocks to Treasury supply propagate internationally not through portfolio rebalancing or interest rates, but through the marginal cost of executing transactions. This is the “safe asset capacity channel” which we formalise in our model. As discussed in Section 3.4, our framework considers the extreme scenario in which stablecoins have fully displaced existing payment systems and capture all settlement activity, both domestic and international, providing an upper bound on the spillovers that Treasury-backed stablecoins can generate.

Exchange rate revaluation channel (novel, amplified by endogenous prices). Since stablecoins are USD-denominated but settle transactions invoiced in various currencies, exchange rate movements create asymmetric effects. Dollar appreciation shrinks the USD value of foreign invoices, reducing stablecoin demand and lowering fees — a positive spillover. Dollar depreciation inflates foreign invoices in USD terms, raising settlement demand and fees — a negative spillover. In our specification, this revaluation channel interacts with endogenous Foreign producer prices: when Foreign productivity rises, Foreign goods become cheaper in FCU terms, compressing invoice values directly and amplifying the positive spillover from dollar appreciation. Home productivity shocks therefore generate larger settlement-demand responses than Foreign shocks of equal magnitude, because the latter are partially self-offsetting through the cost-deflation channel.

The remainder of the paper formalises these channels in a quantitative two-country model, characterises the utilisation threshold that separates the dormant from the active regime, and evaluates the conditions under which foreign CBDC can serve as an effective circuit-breaker.

3 Model environment and timing

We consider two regions: Home (H, interpreted as the United States) and Foreign (F, the rest of the world). Time is discrete, $t = 0, 1, 2, \ldots$. The model is real (flexible prices) with international trade in goods and a single internationally traded financial asset (USD risk-free bond). A USD-denominated stablecoin is introduced as a *settlement technology* for goods transactions.

3.1 Key modelling choices

(A1) **Flexible prices (RBC).** There are no nominal rigidities and no monetary policy rules. Home producer prices are normalised to unity, $P_{HH,t} = 1$ for all t, establishing the US dollar as the model's unit of account. Foreign producer prices are determined endogenously by competitive marginal cost pricing (see Section 3.3).

(A2) **Production with capital and labour.** Each country produces a country-specific intermediate good with Cobb–Douglas technology using predetermined capital and endogenous labour. Labour is chosen by households; the disutility weight is calibrated to match a target steady-state labour level.

(A3) **Armington trade.** In each country, the final good is a CES aggregate of Home and Foreign intermediate goods. Final-good absorption is $X_{j,t} = C_{j,t} + I_{j,t}$.

(A4) **Nominal units and exchange rate.** Home nominal objects are in USD, Foreign nominal objects are in a Foreign currency unit (FCU). The exchange rate is E_t measured as *FCU per USD.* Converting FCU to USD uses division by E_t.

(A5) **International financial market (USD bond + risk premium).** Households trade a one-period USD risk-free bond in zero net supply. A debt-elastic risk premium (in the UIP relation) ensures stationarity of net foreign assets.

(A6) **USD stablecoin settlement with backing scarcity.** Stablecoins are nominal claims denominated in USD that pay zero interest. They are required for settlement of goods purchases in USD. Aggregate settlement demand is backed by short-term US public debt ("bills") through a capacity parameter θ, and a convex *fee* increases with utilisation of backing capacity (see Section 3.4 below).

3.2 Timing within period t

A timing consistent with the implementation in the quantitative model is:

(T1) Productivity, interest-rate, and backing-capacity shocks are realised; prices and the exchange rate adjust.

(T2) Households choose consumption, investment, labour, and the end-of-period USD bond position.

(T3) Goods are traded and *settled* using USD stablecoins. Settlement needs are proportional to nominal invoice values and are grossed up by the endogenously determined stablecoin transaction fee.

(T4) Production takes place; capital is updated for next period; the USD bond pays off.

3.3 Technology and goods

3.3.1 Intermediate-good production

In country $j \in \{H, F\}$, the intermediate good is produced according to

$$Y_{j,t} = A_{j,t}\, K_{j,t-1}^{\alpha}\, N_{j,t}^{1-\alpha}, \tag{1}$$

where $A_{j,t}$ is productivity, $K_{j,t-1}$ is predetermined capital, and $N_{j,t}$ is labour. Capital accumulates as

$$K_{j,t} = (1-\delta)\, K_{j,t-1} + I_{j,t}. \tag{2}$$

3.3.2 Armington final good

In each country, a final good is a CES aggregate of Home and Foreign intermediate goods. For Home:

$$X_{H,t} = \left[(1-\omega_H)^{\frac{1}{\nu}} X_{HH,t}^{\frac{\nu-1}{\nu}} + \omega_H^{\frac{1}{\nu}} X_{FH,t}^{\frac{\nu-1}{\nu}}\right]^{\frac{\nu}{\nu-1}}, \tag{3}$$

and for Foreign:

$$X_{F,t} = \left[(1-\omega_F)^{\frac{1}{\nu}} X_{FF,t}^{\frac{\nu-1}{\nu}} + \omega_F^{\frac{1}{\nu}} X_{HF,t}^{\frac{\nu-1}{\nu}}\right]^{\frac{\nu}{\nu-1}}. \tag{4}$$

Final-good absorption is $X_{j,t} = C_{j,t} + I_{j,t}$.

3.3.3 Prices and exchange-rate mapping

Producer prices. Home intermediate goods are priced in USD and normalised to unity, $P_{HH,t} = 1$ for all t. This establishes the dollar as the model's unit of account and is without loss of generality in a real, flexible- price economy: all terms-of-trade dynamics are absorbed by the exchange rate E_t.

Foreign intermediate goods are priced in FCU and determined endogenously by competitive marginal cost pricing. With a Cobb–Douglas production function and competitive factor markets, the Foreign producer price equals the ratio of the nominal wage to the marginal product of labour:

$$P_{FF,t} = \frac{\bar{W}_F \, N_{F,t}}{(1-\alpha) \, Y_{F,t}}, \tag{5}$$

where $\bar{W}_F$ is a wage normalisation constant calibrated so that $P_{FF,ss} = 1$ in steady state.[3] When Foreign productivity rises, $Y_{F,t}$ increases relative to $N_{F,t}$, so $P_{FF,t}$ falls — Foreign goods become cheaper in FCU terms, which feeds into invoice values and settlement demand through the exchange rate.

Cross-currency prices. With E_t defined as FCU per USD, the bilateral prices are

$$P_{FH,t} = \frac{P_{FF,t}}{E_t} \qquad \text{(Foreign good price in USD)}, \tag{6}$$

$$P_{HF,t} = E_t \, P_{HH,t} = E_t \qquad \text{(Home good price in FCU)}. \tag{7}$$

CES price indices. The final-good price indices are

$$P_{H,t} = \left[(1-\omega_H) \, P_{HH,t}^{1-\nu} + \omega_H \, P_{FH,t}^{1-\nu}\right]^{\frac{1}{1-\nu}}, \qquad P_{F,t} = \left[(1-\omega_F) \, P_{FF,t}^{1-\nu} + \omega_F \, P_{HF,t}^{1-\nu}\right]^{\frac{1}{1-\nu}}. \tag{8}$$

Under the symmetric calibration ($\omega_H = \omega_F$, $E_{ss} = 1$, $P_{HH,ss} = P_{FF,ss} = 1$), both indices satisfy $P_{H,ss} = P_{F,ss} = 1$ in steady state. Out of steady state, $P_{H,t}$ varies with E_t through the import-price component $P_{FH,t} = P_{FF,t}/E_t$, and $P_{F,t}$ varies with both E_t (through $P_{HF,t} = E_t$) and $P_{FF,t}$ (through Foreign productivity). These variations feed into invoice values, settlement demand, and the stablecoin fee.

3.4 Stablecoins as a USD settlement technology

Stablecoins are nominal USD claims paying zero interest, used *within the period* to settle goods purchases. Settlement balances are pinned down by contemporaneous invoice values; the economically relevant wedge is the proportional *fee* paid to execute settlement.

3.4.1 Invoices and required settlement balances

Let $\text{INV}_{H,t}$ and $\text{INV}_{F,t}^{\text{USD}}$ denote the USD value of Home and Foreign absorption, respectively:

$$\text{INV}_{H,t} \equiv P_{H,t} X_{H,t}, \qquad \text{INV}_{F,t}^{\text{USD}} \equiv \frac{P_{F,t} X_{F,t}}{E_t}, \tag{9}$$

---

[3] Specifically, $\bar{W}_F = (1-\alpha) \, Y_{F,ss}/N_{F,ss}$, computed from the symmetric RBC steady state with $r^* = 1/\beta - (1-\delta)$. This ensures that $P_{FF,t}$ fluctuates around unity but is not constrained to equal it at every date.

with E_t in FCU per USD. Required USD stablecoin balances are

$$S_{H,t} = (1 + fee_t)\,\mathrm{INV}_{H,t}, \quad S_{F,t} = (1 + fee_t)\,\mathrm{INV}_{F,t}^{\mathrm{USD}}, \quad S_t \equiv S_{H,t} + S_{F,t}. \tag{10}$$

Settlement must cover the invoice value plus the proportional fee.

Uniform proportional fee. The fee is levied as the same proportion of every Armington variety, so the ratio $(1+fee_t)P_{FH,t}/(1+fee_t)P_{HH,t} = P_{FH,t}$ is independent of the fee. Settlement costs therefore behave like a uniform *ad valorem* sales tax: they reduce households' purchasing power — entering the budget constraints (15)–(16) — but do not distort relative prices in the Armington demand system (3)–(4). The fee operates as a transfer from users to the stablecoin issuer rather than as a trade wedge. The real effects of fee fluctuations are consequently driven by wealth and income channels rather than by expenditure switching; an alternative specification in which the fee falls only on cross-border transactions would add an explicit import-tariff channel that we deliberately abstract from here.

Scope of settlement. Equation (9) applies the stablecoin layer to total final-good absorption in each country, not only to cross-border trade. This is an upper-bound assumption: a stablecoin system confined to international transactions could in principle reach the same utilisation threshold but only at a much larger stablecoin stock relative to Treasury supply. We choose the broader specification because the safe-asset capacity channel is threshold-dependent (Section 6.1): what matters is whether utilisation crosses the convex region of the fee schedule, not the literal share of payments that flow through stablecoins. The parameters B^{short} and θ should be read as calibrated to the combined penetration achieved over the relevant horizon.

3.4.2 Backing constraint and scarcity fee

Aggregate settlement is bounded by available Treasury backing:

$$S_t \le \theta\, B_t^{\mathrm{short}}, \tag{11}$$

where B_t^{short} is the stock of short-dated Treasury instruments available to back stablecoin issuance and $\theta > 0$ is a velocity parameter capturing turnover. Although the GENIUS Act permits a broader set of reserve assets, the empirical evidence (Section 2) shows that profit-maximising issuers concentrate reserves in interest-bearing Treasury instruments at roughly 85% of the total; B_t^{short} therefore captures the marginal binding reserve asset, with θ absorbing both within-period turnover and any dilution from non-Treasury components.

Define utilisation

$$u_t \equiv \frac{S_t}{\theta\, B_t^{\mathrm{short}}} \in [0, 1). \tag{12}$$

We posit a settlement fee increasing and convex in u_t:

$$fee_t = \kappa\, \frac{u_t^p}{1 - u_t}, \qquad \kappa > 0,\ p \ge 1. \tag{13}$$

The reduced form captures, at low utilisation, an almost frictionless settlement layer, and at high utilisation the scarcity premia, balance-sheet costs and repo specialness that issuers face when they must compete for an increasingly limited Treasury collateral pool. The mechanism is consistent with the empirical observation that stablecoin issuance growth tracks Treasury holdings closely (Ferrari Minesso and Siena, 2026; Ahmed and Aldasoro, 2025) and that stablecoin flows have detectable effects on Treasury bill yields.

3.5 Households and international asset market

3.5.1 Preferences

In each country $j \in \{H, F\}$, the representative household maximises

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_{j,t}^{1-\sigma} - 1}{1-\sigma} - \chi \frac{N_{j,t}^{1+\varphi}}{1+\varphi} \right), \tag{14}$$

with $\beta \in (0,1)$, $\sigma > 0$, and $\varphi \geq 0$. Stablecoins do not enter utility.

3.5.2 Budget constraints with settlement fees

Households trade the final good, accumulate capital, and can trade a one-period USD risk-free bond. Stablecoins serve as a payment technology within the period. The associated *fee* is a proportional resource cost paid when executing purchases and therefore enters the budget constraint.

Home (USD). Let $B_{H,t}$ denote Home end-of-period holdings of the USD bond, paying gross return R_t in USD between t and $t{+}1$. The Home stablecoin issuer is owned by Home households: fee revenue collected from Home users is distributed back to Home households as dividends, netting to zero in the consolidated Home budget. The *only* net fee income for Home households is the cross-border transfer $FEE_t^F = fee_t \cdot INV_{F,t}^{USD}$ collected from Foreign users. The Home budget constraint, expressed in USD, is therefore

$$P_{H,t}(C_{H,t} + I_{H,t}) + B_{H,t} = W_{H,t}\, N_{H,t} + R_{K,t}\, K_{H,t-1} + R_t\, B_{H,t-1} + FEE_t^F, \tag{15}$$

where $R_{K,t} = \alpha Y_{H,t}/K_{H,t-1}$ is the competitive rental rate on capital and FEE_t^F is the net stablecoin dividend from Foreign users. The settlement fee on Home's own purchases, $fee_t \cdot INV_{H,t}$, does not appear: it is paid by Home households as consumers and received back in full as shareholders of the Home-based issuer, leaving marginal decisions unaffected. Since $P_{HH,t} = 1$, the CES price index $P_{H,t}$ equals unity in the no-fee steady state and fluctuates around a value close to unity in the true steady state; out of steady state, $P_{H,t}$ varies with E_t through the import-price component $P_{FH,t} = P_{FF,t}/E_t$ of the CES aggregate (8), and $P_{FF,t}$ itself varies endogenously with Foreign productivity through (5).

Foreign (FCU). Let $B_{F,t}$ denote Foreign end-of-period holdings of the USD bond (in USD units). Foreign households do *not* own the stablecoin issuer. The settlement fee is therefore a genuine resource cost that reduces their purchasing power. Writing the Foreign budget in FCU, USD payoffs are converted using E_t. The Foreign budget constraint is

$$(1 + \mathit{fee}_t)\, P_{F,t}(C_{F,t} + I_{F,t}) + E_t\, B_{F,t} = W_{F,t}\, N_{F,t} + R^F_{K,t}\, K_{F,t-1} + E_t\, R_t\, B_{F,t-1}, \tag{16}$$

where $R^F_{K,t} = P_{FF,t}\, \alpha Y_{F,t}/K_{F,t-1}$ is the competitive capital rental rate in FCU. The effective purchase price of Foreign final-good absorption is $(1+\mathit{fee}_t)P_{F,t}$ because the fee applies uniformly to all goods purchases and Foreign receives no offsetting dividend. Note that $(1+\mathit{fee}_t)P_{F,t}X_{F,t} = P_{F,t}X_{F,t} + E_t\, \mathit{fee}_t\, INV^{USD}_{F,t}$, since $INV^{USD}_{F,t} = P_{F,t}X_{F,t}/E_t$, so the two equivalent representations of the Foreign budget are related by this identity.

3.5.3 Labour supply

Competitive factor markets and cost minimisation yield nominal wages $W_{H,t} = (1-\alpha)Y_{H,t}/N_{H,t}$ (since $P_{HH,t} = 1$) and $W_{F,t} = P_{FF,t}(1-\alpha)Y_{F,t}/N_{F,t}$. The labour supply conditions follow from each household's budget constraint.

For Home households, the settlement fee nets to zero in the consolidated budget (equation (15)), so the effective purchase price of consumption is simply $P_{H,t}$. The Home labour FOC is standard:

$$\chi\, N^{\varphi}_{H,t} = C^{-\sigma}_{H,t}\, \frac{(1-\alpha)\, Y_{H,t}}{N_{H,t}\, P_{H,t}}. \tag{17}$$

For Foreign households, the effective purchase price of consumption is $(1 + \mathit{fee}_t)P_{F,t}$ (equation (16)), which enters the real consumption wage in the denominator. The Foreign labour FOC is:

$$\chi\, N^{\varphi}_{F,t} = C^{-\sigma}_{F,t}\, \frac{P_{FF,t}}{P_{F,t}}\, \frac{(1-\alpha)\, Y_{F,t}}{N_{F,t}\,(1 + \mathit{fee}_t)}. \tag{18}$$

The ratio $P_{FF,t}/P_{F,t}$ converts the FCU product wage into real consumption units; the additional $(1 + \mathit{fee}_t)$ in the denominator reflects the fee's role as a tax on Foreign household purchasing power. Both ratios equal unity in steady state but fluctuate along the transition, opening channels from Foreign cost conditions and from the settlement fee to Foreign labour supply that are absent from the Home equation.

3.5.4 Capital Euler conditions

The Euler equations for capital follow from the household optimality conditions applied to each country's budget constraint. The key result, derived from the Lagrangian of each household's problem, is that the shadow value of installed capital equals the marginal utility of consumption regardless of the fee, because investment and consumption face the same proportional cost within a period and the fee therefore cancels in the investment-versus-consumption comparison. The fee enters the Euler only through the return on capital realised next period, because rental income is revenue that must be spent through the stablecoin system, whereas the undepreciated capital $(1-\delta)K_{j,t}$ rolls forward as a physical quantity free of any fee distortion.

For **Home**, the fee on domestic transactions nets to zero after dividends, so the effective purchase price of consumption is $P_{H,t}$ and the Euler takes the standard RBC form:

$$C_{H,t}^{-\sigma} = \beta\, \mathbb{E}_t \left[C_{H,t+1}^{-\sigma} \left(\frac{\alpha\, Y_{H,t+1}}{K_{H,t}\, P_{H,t+1}} + 1 - \delta \right) \right]. \tag{19}$$

The presence of $1/P_{H,t+1}$ rather than $1/P_{HH,t+1} = 1$ reflects the CES basket price: since $P_{HH,t} = 1$, the real return to capital is the marginal product deflated by the consumer price index, not the producer price.

For **Foreign**, the effective purchase price of consumption is $(1 + \mathit{fee}_t)P_{F,t}$. The Lagrange multiplier on the Foreign budget constraint is therefore $\lambda_{F,t} = C_{F,t}^{-\sigma}/[(1 + \mathit{fee}_t)P_{F,t}]$, which introduces $(1 + \mathit{fee}_{t+1})$ into the denominator of the MPK term at $t + 1$ when converting rental income into utility:

$$C_{F,t}^{-\sigma} = \beta\, \mathbb{E}_t \left[C_{F,t+1}^{-\sigma} \left(\frac{P_{FF,t+1}}{P_{F,t+1}} \frac{\alpha\, Y_{F,t+1}}{K_{F,t}\, (1 + \mathit{fee}_{t+1})} + 1 - \delta \right) \right]. \tag{20}$$

Equation (20) shows that a higher expected settlement fee at $t + 1$ reduces the effective real return on Foreign capital, providing an additional incentive for Foreign households to bring consumption forward and reduce current investment. This intertemporal distortion channel is absent from Home's Euler (19), generating an asymmetry in investment dynamics between the two countries.

The two Euler equations imply different steady-state capital stocks. From (19), Home capital satisfies the standard condition $K_{H,ss} = \alpha Y_{H,ss}/(r^* P_{H,ss})$ where $r^* = 1/\beta - (1 - \delta)$. From (20), Foreign capital satisfies $K_{F,ss} = \alpha Y_{F,ss}/[r^*(1 + \mathit{fee}_{ss})P_{F,ss}]$. Since $\mathit{fee}_{ss} > 0$, we have $K_{F,ss} < K_{H,ss}$ at any symmetric calibration, implying a slightly asymmetric steady state in which Foreign is a net borrower from Home through the stablecoin fee mechanism. The asymmetry is of order $\mathit{fee}_{ss} \approx 0.1\%$ and is reproduced exactly by the numerical steady-state solver.

3.5.5 UIP with debt-elastic risk premium

Let i_t^{US} and i_t^F denote the Home and Foreign nominal interest rates (exogenous AR(1) processes in the quantitative model). The exchange rate satisfies an uncovered interest parity condition with a debt-elastic risk premium:

$$(1 + i_t^{US}) = (1 + i_t^F)\, \frac{E_t}{E_{t+1}}\, \exp\{\phi_B(BF_t - \bar{B})\}, \tag{21}$$

where BF_t is Foreign net holdings of the USD bond (in USD units), $\phi_B > 0$, and $\bar{B}$ is the steady-state level.

Equation (21) is the consolidated no-arbitrage condition implied by the two countries' intertemporal optimality conditions for the USD bond. The Home household's FOC for $B_{H,t}$ takes the standard form $C_{H,t}^{-\sigma}/P_{H,t} = \beta R_t\, \mathbb{E}_t[C_{H,t+1}^{-\sigma}/P_{H,t+1}]$ with $R_t = 1 + i_t^{US}$. The Foreign household holds the same USD bond, but prices its position in FCU at the exchange rate E_t and perceives an effective return $R_t \exp\{-\phi_B(BF_t - \bar{B})\}$, where the debt-elastic wedge reduces the marginal value of accumulating USD assets. Combining the Foreign household's FOC for

$B_{F,t}$ with the analogous (implicit) FOC for an FCU-denominated bond paying $1 + i_t^F$ yields equation (21); the settlement-fee terms $(1 + \mathit{fee}_t)P_{F,t}$ enter both Foreign Eulers symmetrically and cancel in the no-arbitrage ratio to first order, so the UIP relation does not carry the settlement wedge directly. The premium $\phi_B(BF_t - \bar{B})$ ensures a unique stationary distribution of Foreign USD assets in this open-economy formulation, as in Schmitt-Grohé and Uribe (2003).

3.5.6 Foreign net foreign assets (USD)

Define the Foreign trade balance (in USD) as

$$TBF_t^{USD} = P_{FH,t}\, X_{FH,t} - P_{HH,t}\, X_{HF,t} = P_{FH,t}\, X_{FH,t} - X_{HF,t}. \tag{22}$$

Foreign net foreign assets in the USD bond evolve as

$$BF_t = (1 + i_{t-1}^{US})\, BF_{t-1} + TBF_t^{USD} - FEE_t^F, \tag{23}$$

so that fee revenue collected on Foreign invoices reduces Foreign USD net assets one-for-one.

Interpretation of fee transfers. The term FEE_t^F in equation (23) represents the net fee payment from Foreign to Home households, measured in USD. This transfer arises because the stablecoin issuer is a US-based firm owned by Home households. When Foreign agents settle USD-denominated transactions, they pay a fee to the issuer:

$$FEE_t^F = \mathit{fee}_t \cdot INV_{F,t}^{USD}.$$

These revenues accrue to the Home-based issuer as operating profits. Home households also pay fees on their own transactions:

$$FEE_t^H = \mathit{fee}_t \cdot INV_{H,t}.$$

The dividend received by Home households is therefore

$$\mathit{Dividends}_t^{\mathit{stablecoin}} = FEE_t^F + FEE_t^H = \mathit{fee}_t \cdot \big(INV_{F,t}^{USD} + INV_{H,t}\big).$$

Only FEE_t^F appears in the NFA accumulation equation (23), because only the cross-border fee transfer affects the international balance of payments. Home's fee payments to its own issuer are purely domestic transfers.

This ownership structure has a crucial implication for household optimality conditions. Because Home households pay FEE_t^H as consumers and receive it back as shareholders, the fee is a *lump-sum circular transfer* at the household level: it does not affect the marginal cost of any decision. The effective purchase price of Home absorption is therefore simply $P_{H,t}$, and the settlement fee does not appear in Home's labour supply condition or capital Euler equation. Foreign households, by contrast, face a genuine resource cost: they pay FEE_t^F with no offsetting dividend, so the fee raises their effective cost of consumption to $(1 + \mathit{fee}_t)P_{F,t}$ and distorts both their labour supply and their intertemporal saving decisions.

A digital exorbitant privilege. The asymmetric structure of the cross-border fee transfer in equation (23) constitutes a digital counterpart to the dollar's classical *exorbitant privilege* of Gourinchas and Rey (2007); Gourinchas et al. (2019). In the original formulation, the US earns a permanent return differential on its external balance sheet — a higher yield on foreign assets than the rate paid on foreign liabilities — which generates a steady-state income flow from the rest of the world even when net positions are small. In our model the classical channel is absent by construction: both countries hold the same USD bond at the same gross return R_t, the bond-market clearing $B_{H,t} + BF_t = 0$ together with symmetric productivity implies $BF_{ss} \approx 0$ at the steady state, and there is no yield differential to exploit. Despite this, Foreign households pay a perpetual flow $FEE_t^F = fee_t \cdot INV_{F,t}^{USD}$ to Home households as compensation for using the USD settlement infrastructure that Home owns. The mechanism is structurally identical to the classical privilege — a one-way international income transfer arising from Home's ownership of a globally-demanded service — but operates through fees on payment infrastructure rather than through asset-return differentials. The privilege resides in ownership of the dominant settlement layer rather than in the composition of the external balance sheet, so it remains operative even in steady states where the classical return-differential channel is muted.

3.6 Equilibrium

An equilibrium is a set of sequences for quantities, prices, and asset positions

$$\left\{C_{j,t}, I_{j,t}, K_{j,t}, N_{j,t}, X_{ab,t}, Y_{j,t}, P_{H,t}, P_{F,t}, P_{FF,t}, E_t, S_{H,t}, S_{F,t}, S_t, u_t, fee_t, BF_t\right\}_{t\geq 0}$$

such that: (i) households optimise given prices; (ii) firms satisfy production and goods-market clearing; (iii) Armington demands and price indices hold; (iv) Foreign producer prices satisfy (5); (v) the exchange rate satisfies (21); (vi) settlement demand and fees satisfy (9)–(13); and (vii) Foreign net foreign assets follow (23).

3.6.1 Market clearing

Intermediate goods.

$$Y_{H,t} = X_{HH,t} + X_{HF,t}, \qquad Y_{F,t} = X_{FF,t} + X_{FH,t}. \tag{24}$$

Final-good absorption.

$$X_{H,t} = C_{H,t} + I_{H,t}, \qquad X_{F,t} = C_{F,t} + I_{F,t}. \tag{25}$$

Bond market.

$$B_{H,t} + BF_t = 0. \tag{26}$$

Stablecoin settlement (global).

$$S_t = S_{H,t} + S_{F,t}, \tag{27}$$

with $S_{H,t}$ and $S_{F,t}$ pinned down by invoice values and the utilisation-based fee.

3.7 Flex prices and the dominant currency paradigm

The model is cast in real terms with flexible prices, an RBC structure that abstracts from nominal rigidities and from independent monetary policy. Three remarks on this choice and on its relation to the dominant currency paradigm (DCP) of Gopinath et al. (2020) are in order.

First, under flexible prices the distinction between producer currency pricing, local currency pricing, and dominant currency pricing collapses: all prices adjust instantaneously to clear markets, and the law of one price holds across markets net of transport and settlement costs. The flex-price equilibrium described above is therefore consistent with all three pricing conventions at the steady state. The normalisation $P_{HH,t} = 1$ and the endogeneity of the Foreign producer price $P_{FF,t}$ are a modelling convenience, not a substantive assumption about invoicing currency.

Second, the safe asset capacity channel operates through the USD-invoice volume $INV_{F,t}^{USD}$ and the resulting stablecoin fee, not through nominal price stickiness. The channel is therefore conceptually orthogonal to the DCP mechanism: DCP explains *why* so much international trade is invoiced in USD; our model takes the invoicing structure as given through the parameters ω_H^{USD} and ω_F^{USD} and asks what happens when the *settlement layer* for those USD-invoiced transactions runs into Treasury-backing capacity constraints. The two literatures are complementary — the DCP literature explains why USD invoicing is dominant, while ours explains how that dominance is propagated to real outcomes through endogenous settlement frictions.

Third, what would reasonably happen under a sticky-price assumption? Under DCP, the cross-border USD-invoiced share of trade prices is sticky in USD rather than in the importer's currency. An exchange rate depreciation does not pass through to Foreign import prices in the short run, so the real burden of those imports on Foreign households increases. Combined with the stablecoin settlement wedge $\tau_{F,t}$, this would generate a compounded purchasing-power loss for Foreign agents: both the sticky USD price level and the USD-settlement wedge act in the same direction, amplifying the spillover. By contrast, local-currency pricing would partly insulate Foreign households from USD-side cost shocks, attenuating the channel. Our flex-price results therefore plausibly *understate* the spillover magnitude in a DCP world and *overstate* it in a purely local-currency-pricing world; the channel itself — driven by safe-asset scarcity on the stablecoin issuer's balance sheet rather than by nominal frictions — should remain qualitatively operative in either case. A full DCP–NK extension, in which monetary policy and exchange-rate pass-through interact with the fee channel, is left for future work.

4 Introducing a Foreign Currency CBDC

This section extends the baseline model to explore how the introduction of a costless foreign Central Bank Digital Currency (CBDC) affects the international monetary system and the transmission mechanisms identified in the previous section. The key question is: *if the Foreign central bank issues a digital currency that can be used for settlement at zero cost, would USD stablecoins retain their dominant role, and how would the "safe asset capacity channel" be affected?*

4.0.1 Motivation and context

The rapid development of CBDC projects worldwide — including the digital euro, digital yuan, and dozens of pilot programmes in emerging economies — reflects concerns about private stablecoin dominance and the desire to maintain monetary sovereignty in an increasingly tokenised payment landscape. A key feature distinguishing CBDCs from private stablecoins is their status as *public money* issued directly by central banks, which in principle can be supplied elastically without reliance on scarce collateral such as Treasury bills.

If a major foreign central bank were to issue a digital currency that is: (i) *denominated in the foreign currency* (FCU), providing a natural hedge for non-US residents; (ii) *costless to use for settlement*, avoiding the transaction fees embedded in private stablecoin systems; and (iii) *elastically supplied*, without capacity constraints tied to safe asset availability; then such a CBDC could in principle compete directly with USD stablecoins for settlement market share.

Our extension introduces precisely this scenario. We model a Foreign CBDC that households in both countries can use to settle transactions denominated in foreign currency, and allow agents to choose endogenously between USD stablecoins (with positive fees and capacity constraints) and the Foreign CBDC (costless and unconstrained). This setup enables us to address various issues, including the coexistence conditions under which both settlement systems remain active, the competitive discipline imposed by CBDC availability on USD stablecoin fees, the stabilisation role of CBDCs in absorbing settlement demand during periods of USD liquidity stress, and the degree of decoupling that Foreign can achieve from US safe asset supply shocks.

4.1 Model extension: dual settlement system

We modify the baseline framework to allow transactions to be settled using either USD stablecoins or Foreign CBDC, depending on the currency of invoicing and households' cost-benefit trade-offs.

4.1.1 Currency-of-invoicing and mandatory settlement shares

A key empirical regularity in international trade is *dominant currency pricing*: a disproportionate share of global trade is invoiced in US dollars even when neither the exporter nor the importer is located in the United States (Maggiori et al., 2020; Gopinath and Stein, 2021). This creates a natural segmentation in settlement demand: transactions invoiced in dollars *must* be settled in dollar-denominated instruments (USD stablecoins in our model), while transactions invoiced in foreign currency *can* be settled in Foreign CBDC.

We introduce two parameters to capture this invoicing structure:

$$\omega_H^{\text{USD}} \in [0,1] : \text{ share of Home absorption invoiced in USD (mandatory stablecoin use)}, \tag{28}$$

$$\omega_F^{\text{USD}} \in [0,1] : \text{ share of Foreign absorption invoiced in USD (mandatory stablecoin use).} \tag{29}$$

Calibration: scope matters. A subtle but important point governs the calibration of these parameters. The classic Gopinath and Stein (2021) shares of dollar invoicing in international trade — approximately 85% for the United States and 40–50% for the rest of the world — refer to *cross-border trade flows* only. In our model, however, the invoice values $INV_{j,t} = P_{j,t}\, X_{j,t}$ defined in equation (9) cover *total final-good absorption* $X_j = C_j + I_j$, including domestic transactions. Applying the cross-border-trade shares directly to total absorption would dramatically overstate USD invoicing for Foreign, since Foreign domestic transactions are predominantly invoiced in the Foreign currency unit (FCU), not USD.

To calibrate ω_j^{USD} consistently with the total-absorption scope of the model, we weight the cross-border-trade USD share by the import share $\omega_j = 0.25$ that, in the Armington structure of Section 3.3.2, fixes the fraction of absorption that crosses borders:

$$\omega_F^{\text{USD}} \approx \underbrace{0.75 \times 0}_{\text{domestic, FCU}} + \underbrace{0.25 \times 0.45}_{\text{cross-border, 45\% USD}} \approx 0.10, \tag{30}$$

$$\omega_H^{\text{USD}} \approx \underbrace{0.75 \times 1.00}_{\text{domestic, USD}} + \underbrace{0.25 \times 0.80}_{\text{cross-border, 80\% USD}} \approx 0.95. \tag{31}$$

The Home value is near unity because home-currency pricing means that even US cross-border transactions are predominantly invoiced in dollars; we conservatively set the cross-border USD share at 80%, slightly below the ≈93% reported by Gopinath and Stein (2021), which only understates Home dollar dominance. The Foreign value is much smaller than the headline 45% because the dollar's invoicing role in the rest of the world is concentrated in cross-border trade rather than in everyday domestic transactions.

We adopt these as our baseline calibration:

$$\omega_H^{\text{USD}} = 0.95, \tag{32}$$

$$\omega_F^{\text{USD}} = 0.10. \tag{33}$$

The precise value of ω_H^{USD} is immaterial for our results: any value in $[0.95, 1.0]$ yields essentially identical outcomes, since Home's discretionary margin is negligible and Home settles almost entirely in USD in equilibrium regardless (the baseline CBDC steady state has $share_H^{\text{USD}} = 0.98$). The substantive recalibration relative to a naive reading of the invoicing data is the Foreign floor ω_F^{USD}.

These mandatory invoicing floors are deliberately low — in particular for Foreign, where $1 - \omega_F^{\text{USD}} = 0.90$ of absorption is *discretionary* and can in principle be diverted to a costless foreign CBDC. Whether stablecoins retain settlement dominance in the CBDC extension is therefore not an artefact of high mandatory floors imposed by calibration: it reflects the equilibrium interaction of cost differentials and network effects on the discretionary share, which we analyse quantitatively in Section 6.1.

The *complementary shares* $(1 - \omega_j^{\text{USD}})$ represent transactions invoiced in foreign currency, for which agents have discretion to choose between USD stablecoins and Foreign CBDC based on relative costs and network considerations.

4.1.2 Invoice values by currency

Total invoice values in each country are split by currency of denomination. For Home:

$$INV_H^{\text{USD}} = P_{H,t}\,X_{H,t}, \tag{34}$$

$$INV_H^{\text{FCU}} = P_{H,t}\,X_{H,t} \cdot E_t, \tag{35}$$

and for Foreign:

$$INV_F^{\text{USD}} = \frac{P_{F,t}\,X_{F,t}}{E_t}, \tag{36}$$

$$INV_F^{\text{FCU}} = P_{F,t}\,X_{F,t}. \tag{37}$$

The exchange rate E_t is central to the revaluation effects: when the dollar appreciates ($E_t \uparrow$), the USD value of Foreign invoices INV_F^{USD} falls mechanically, reducing Foreign demand for USD stablecoins. At the same time, $P_{F,t}$ responds to Foreign productivity through $P_{FF,t}$, generating an additional *cost-deflation channel*: a positive Foreign productivity shock lowers $P_{FF,t}$, compresses $P_{F,t}$, and directly reduces INV_F^{USD}, partially offsetting the invoice inflation that would otherwise arise from increased Foreign absorption.

4.1.3 Endogenous settlement shares with network effects

Households in each country choose how much of their *discretionary* transaction volume to settle using USD stablecoins versus Foreign CBDC. In our model this choice reflects:

- **Cost differential:** USD stablecoins charge an endogenous $fee_t > 0$, while CBDC settlement is costless.
- **Network externalities:** Settlement systems exhibit natural network effects — an instrument is more valuable when more counterparties use it. Higher aggregate USD stablecoin usage increases the convenience and acceptance of dollar settlement, creating strategic complementarity in settlement choices.

We model the *settlement share* using a logistic function that nests both cost sensitivity and network effects:

$$share_{H,t}^{\text{USD}} = \omega_H^{\text{USD}} + \frac{1 - \omega_H^{\text{USD}}}{1 + \exp\left[-\eta \cdot \dfrac{S_t}{S_t + CBDC_t^{\text{USD}}} + \psi \cdot fee_t\right]}, \tag{38}$$

$$share_{F,t}^{\text{USD}} = \omega_F^{\text{USD}} + \frac{1 - \omega_F^{\text{USD}}}{1 + \exp\left[-\eta \cdot \dfrac{S_t}{S_t + CBDC_t^{\text{USD}}} + \psi \cdot fee_t\right]}, \tag{39}$$

where $CBDC_t^{\text{USD}} = CBDC_t/E_t$ is CBDC usage expressed in USD, $\eta > 0$ governs the strength of network externalities, and $\psi > 0$ is a rescaling parameter that converts the fee (expressed as a decimal) into units comparable to the network term (which lies in $[0, 1]$). We set $\psi = 100$ in the quantitative implementation, so that $\psi \cdot fee_t$ is the fee expressed in percentage points.[4]

Equations (38)–(39) say that each country's USD settlement share is bounded below by the mandatory invoicing share (ω_j^{USD}) and can rise above it depending on the balance between network benefits η (which favour USD when its market share is high) and cost penalties (which favour CBDC when fee_t is high). The logistic form ensures $share_{j,t}^{\mathrm{USD}} \in [\omega_j^{\mathrm{USD}}, 1]$ and generates smooth, continuous adjustment.

Figure 1 illustrates the dual role of network effects and fee penalties in determining Foreign USD settlement choices.

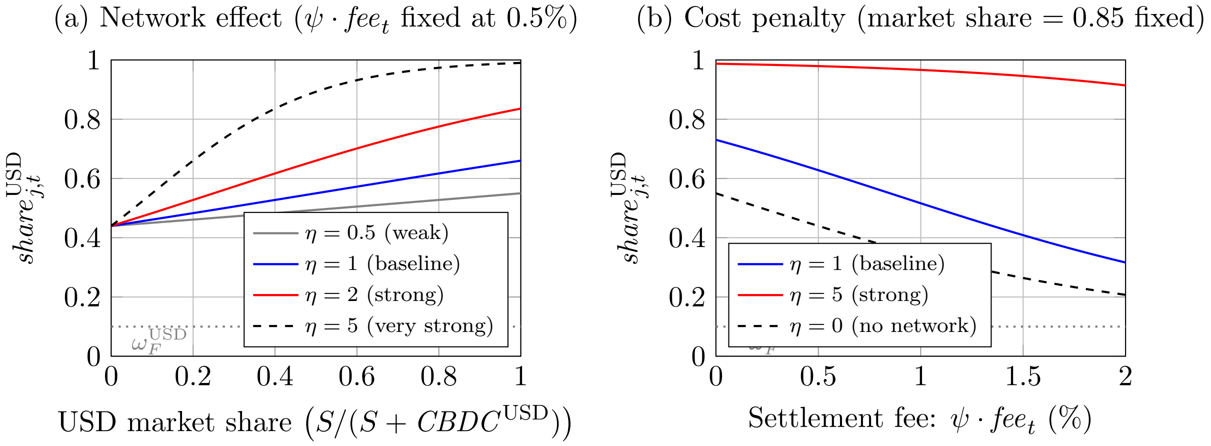

(a) Network effect ($\psi \cdot fee_t$ fixed at 0.5%)

(b) Cost penalty (market share = 0.85 fixed)

Figure 1: Endogenous Foreign USD settlement share as a function of network effects and fees, under the calibration $\omega_F^{\mathrm{USD}} = 0.10$ and $\psi = 100$. Panel (a) shows how the settlement share responds to the prevailing USD market share $S/(S + CBDC^{\mathrm{USD}})$ for different network-effect strengths η, holding the fee fixed at $\psi \cdot fee_t = 0.5\%$. At the baseline strength $\eta = 1$, network effects lift the Foreign USD share materially above the mandatory floor as the USD market share rises (reaching ≈ 0.66 when USD market share equals 1), but do not fully tip the system toward $share_F^{\mathrm{USD}} = 1$. Stronger network effects ($\eta \geq 2$) generate the sharp tipping dynamics associated with self-reinforcing settlement-system dominance. Panel (b) illustrates the cost-penalty channel: holding the USD market share fixed at 0.85 (its baseline-calibration steady-state value), higher stablecoin fees push the discretionary settlement share toward Foreign CBDC. At $\eta = 1$ the share is highly fee-elastic; at $\eta = 5$ network lock-in dominates and the fee penalty barely affects allocation. The case $\eta = 0$ isolates the pure cost-penalty effect with no network reinforcement. The horizontal dotted line marks the mandatory invoicing floor $\omega_F^{\mathrm{USD}} = 10\%$ below which the share cannot fall.

---

[4] At the baseline calibration with CBDC, $fee_{ss} \approx 0.00077$ (about 31 basis points annualised), so $\psi \cdot fee_{ss} \approx 0.077$. The rescaling by $\psi = 100$ raises the fee penalty from a magnitude of order 10^{-3} to a magnitude of order 10^{-1}, making it economically detectable relative to the network term $S_{ss}/(S_{ss} + CBDC_{ss}^{\mathrm{USD}}) \approx 0.85$. The network benefit still dominates the fee penalty at steady state in the baseline regime; the fee becomes a meaningfully binding discipline on settlement choice in the stressed and extreme regimes explored in Section 6.1.

4.1.4 USD stablecoin settlement

Settlement demand for USD stablecoins is determined by the portion of invoices settled in USD, grossed up by the endogenous fee:

$$\begin{aligned} S_{H,t} &= (1 + \mathit{fee}_t) \cdot \mathit{share}^{\text{USD}}_{H,t} \cdot \mathit{INV}^{\text{USD}}_{H,t}, &(40)\\ S_{F,t} &= (1 + \mathit{fee}_t) \cdot \mathit{share}^{\text{USD}}_{F,t} \cdot \mathit{INV}^{\text{USD}}_{F,t}, &(41)\\ S_t &= S_{H,t} + S_{F,t}. &(42) \end{aligned}$$

The utilisation of USD stablecoin backing capacity and the convex fee schedule remain as in the baseline:

$$\begin{aligned} u_t &= \frac{S_t}{\theta \cdot B^{\text{short}}_t}, &(43)\\ \mathit{fee}_t &= \kappa \cdot \frac{u^p_t}{1 - u_t}. &(44) \end{aligned}$$

The aggregate USD stablecoin demand S_t now depends on the endogenous shares $\mathit{share}^{\text{USD}}_{j,t}$, which fall when the fee rises. This introduces a *competitive discipline mechanism*: as utilisation approaches capacity and the fee rises, households shift to CBDC, reducing S_t and moderating the fee increase.

4.1.5 Foreign CBDC settlement

The Foreign central bank issues a digital currency denominated in FCU that is available costlessly for settlement:

$$\begin{aligned} \mathit{CBDC}_{H,t} &= \left(1 - \mathit{share}^{\text{USD}}_{H,t}\right) \cdot \mathit{INV}^{\text{FCU}}_{H,t}, &(45)\\ \mathit{CBDC}_{F,t} &= \left(1 - \mathit{share}^{\text{USD}}_{F,t}\right) \cdot \mathit{INV}^{\text{FCU}}_{F,t}, &(46)\\ \mathit{CBDC}_t &= \mathit{CBDC}_{H,t} + \mathit{CBDC}_{F,t}. &(47) \end{aligned}$$

The Foreign central bank accommodates CBDC demand elastically, without reserve backing constraints.

4.1.6 Why would anyone hold USD stablecoins?

(i) Mandatory invoicing in USD. A fraction ω^{USD}_j of transactions is invoiced in dollars due to contractual and legal reasons (commodity markets, international debt contracts), established trade relationships with dollar pricing, and the desire to hedge against local currency depreciation. These transactions cannot be settled in FCU-denominated CBDC. Even if the Foreign CBDC were universally superior, this invoicing inertia guarantees a baseline demand for USD stablecoins of at least:

$$S^{\min}_t = (1 + \mathit{fee}_t)\Big(\omega^{\text{USD}}_H \cdot \mathit{INV}^{\text{USD}}_{H,t} + \omega^{\text{USD}}_F \cdot \mathit{INV}^{\text{USD}}_{F,t}\Big). \tag{48}$$

(ii) Network externalities and coordination benefits. Beyond mandatory usage, the logistic share equations (38)–(39) embed a *strategic complementarity*: when many agents use USD settlement, it becomes more valuable for additional agents to adopt it as well, because counterparties are more likely to accept USD-denominated settlement, liquidity and market depth are higher in the dominant system, and integration with existing financial infrastructure is stronger for the incumbent system. Our calibration with $\eta = 1.0$ ensures this effect is present but not so strong as to create multiple equilibria; the fee still plays a disciplining role.

4.1.7 Effective settlement cost under CBDC

In the baseline every transaction uses USD stablecoins, so both countries face the same proportional fee on absorption. Under the CBDC extension, the effective settlement cost per unit of absorption is the share-weighted fee

$$\tau_{j,t} = \mathit{share}_{j,t}^{\mathrm{USD}} \cdot \mathit{fee}_t, \tag{49}$$

which households can reduce by routing transactions through the costless Foreign CBDC on the discretionary margin. In equilibrium $\mathit{share}_{H,t}^{\mathrm{USD}} \approx 0.98$ while $\mathit{share}_{F,t}^{\mathrm{USD}} \approx 0.72$, so the gross fee burden is larger for Home in accounting terms. The *economic incidence* runs the other way: the fee revenue collected on Home users returns to Home households as dividends from the US-based issuer, making $\tau_{H,t}$ a circular transfer that drops out of every Home optimality condition. For Foreign households, by contrast, $\tau_{F,t}$ is a genuine resource cost with no offsetting dividend — it raises the effective price of Foreign absorption to $(1 + \tau_{F,t})P_{F,t}$ and enters both the Foreign labour FOC and the Foreign capital Euler with the standard intratemporal and intertemporal distortions analysed in Section 3.5. Because $\tau_{F,t}$ is small in steady state ($\tau_{F,ss} \approx \mathit{share}_{F,ss}^{USD} \cdot \mathit{fee}_{ss} \approx 0.72 \times 0.00077 \approx 0.00055$, i.e. about 5 bp per quarter), χ_F is calibrated essentially symmetrically to χ_H at $N_{F,ss} = 0.33$. The formal FOCs are reported in Section 4.2 below.

Two features of equation (49) are worth flagging. First, the fee applies as a uniform markup to total absorption $P_{j,t}X_{j,t}$ and therefore does not distort within-country Armington demand: relative prices of domestic versus imported intermediates are unaffected, so the channel operates through wealth and income rather than expenditure switching. Second, differentiating (49), $\partial\tau_{F,t}/\partial \mathit{fee}_t = \mathit{share}_{F,t}^{\mathrm{USD}} + \mathit{fee}_t \cdot \partial \mathit{share}_{F,t}^{\mathrm{USD}}/\partial \mathit{fee}_t \approx 0.72$ at the calibrated CBDC steady state: a 10-basis-point fee shock raises Foreign's effective settlement cost by roughly 7.2 basis points, while Home households are insulated by the dividend offset. This asymmetric incidence is the microfoundation for the cross-country macroeconomic results in Section 6.1.

4.2 Household budget constraints and first-order conditions

Under the dual settlement system, only USD stablecoin transactions incur a fee; CBDC settlement is costless. Define the gross fee revenues collected from Home and Foreign users by $\mathit{FEE}_t^j \equiv \mathit{fee}_t \cdot \mathit{share}_{j,t}^{USD} \cdot \mathit{INV}_{j,t}^{USD}$, with $j \in \{H, F\}$.

Home (USD). Home households own the stablecoin issuer, so FEE_t^H is rebated to them as a dividend and nets to zero in the consolidated Home budget. The only net fee inflow to Home is the cross-border transfer FEE_t^F, and the Home budget constraint reads

$$P_{H,t}\, X_{H,t} + B_{H,t} = W_{H,t}\, N_{H,t} + R_{K,t}\, K_{H,t-1} + R_t\, B_{H,t-1} + FEE_t^F, \tag{50}$$

with $R_{K,t} = \alpha Y_{H,t}/K_{H,t-1}$ and $X_{H,t} = C_{H,t} + I_{H,t}$. The effective purchase price of Home absorption is $P_{H,t}$, since the fee paid as consumer is exactly offset by the dividend received as shareholder.

Foreign (FCU). Foreign households do not own the issuer, so FEE_t^F is a genuine resource cost. Writing the Foreign budget in FCU with USD payoffs converted at E_t:

$$(1 + \tau_{F,t})\, P_{F,t}\, X_{F,t} + E_t\, B_{F,t} = W_{F,t}\, N_{F,t} + R_{K,t}^F\, K_{F,t-1} + E_t\, R_t\, B_{F,t-1}, \tag{51}$$

with $\tau_{F,t} = share_{F,t}^{\mathrm{USD}} \cdot fee_t$, $X_{F,t} = C_{F,t} + I_{F,t}$, and $R_{K,t}^F = P_{FF,t}\, \alpha Y_{F,t}/K_{F,t-1}$.

First-order conditions. The Home FOCs are identical to the baseline (eqs. (17) and (19)) because the Home wedge $\tau_{H,t}$ does not enter Home optimality conditions. For Foreign, the labour FOC and the capital Euler replace fee_t with $\tau_{F,t}$:

$$\chi_F\, N_{F,t}^{\varphi} = C_{F,t}^{-\sigma}\, \frac{P_{FF,t}}{P_{F,t}}\, \frac{(1-\alpha)\, Y_{F,t}}{N_{F,t}\,(1 + \tau_{F,t})}, \tag{52}$$

$$C_{F,t}^{-\sigma} = \beta\, \mathbb{E}_t \left[C_{F,t+1}^{-\sigma} \left(\frac{P_{FF,t+1}}{P_{F,t+1}}\, \frac{\alpha\, Y_{F,t+1}}{K_{F,t}\,(1 + \tau_{F,t+1})} + 1 - \delta \right) \right]. \tag{53}$$

The UIP condition (21) carries through unchanged: the substitution $fee_t \mapsto \tau_{F,t}$ appears symmetrically in the Foreign Eulers for the USD and the implicit FCU bond and cancels in the no-arbitrage ratio to first order. The settlement-share variables $share_{j,t}^{\mathrm{USD}}$ that define $\tau_{F,t}$ are pinned down by the logistic specifications (38)–(39) rather than by explicit optimisation, consistent with the empirical regularity that currency choice in international trade is driven by coordination forces and contractual conventions (Gopinath and Stein, 2021).

5 Calibration and simulation

5.1 Baseline parametrization

Table 2 summarizes the parameter values used in both the baseline model (stablecoins only) and the CBDC extension. The calibration follows standard practices in international RBC literature while introducing novel parameters specific to the stablecoin settlement mechanism.

Preferences and technology: The discount factor $\beta = 0.99$ implies a steady-state quarterly real interest rate of approximately 1%, consistent with observed safe rates. The coefficient of relative risk aversion $\sigma = 2$ is standard in macro models. We set the Frisch elasticity $1/\varphi = 1$, implying moderate labor supply responsiveness. The capital share $\alpha = 0.33$ and depreciation rate $\delta = 0.025$ match US data.

Table 2: Parameter calibration

Parameter	**Symbol**	**Value**	**Target / Source**
Preferences and technology			
Discount factor	β	0.99	Annual real rate $\approx 1\%$
Risk aversion	σ	2.0	Standard macro
Frisch elasticity inv.	φ	1.0	Labor supply lit.
Capital share	α	0.33	US data
Depreciation rate	δ	0.025	Quarterly, standard
Trade structure			
CES elasticity	ν	1.5	Trade elasticity lit.
Home import share	ω_H	0.25	Import/GDP $\approx 25\%$
Foreign import share	ω_F	0.25	Symmetric baseline
Stablecoin settlement			
Velocity parameter	θ	0.5	Turnover
Fee scale	κ	0.002	Targets baseline regime fee (50 bp ann.)
Fee convexity	p	2.0	Moderate congestion
Numerical tolerance	ε	10^{-6}	Numerical stability
External adjustment			
Risk premium elast.	ϕ_B	0.05	NFA stationarity
NFA target	$\bar{B}$	0.0	Endogenous (SS)
AR(1) persistence			
Home productivity	ρ_{A_H}	0.95	Standard RBC
Foreign productivity	ρ_{A_F}	0.95	Standard RBC
US interest rate	$\rho_{i^{US}}$	0.90	Policy inertia
Foreign interest rate	ρ_{i^F}	0.90	Policy inertia
Backing capacity	$\rho_{B^{short}}$	0.90	Treasury persistence
Steady-state levels			
US interest rate	i_{ss}^{US}	0.01	1% quarterly
Foreign interest rate	i_{ss}^{F}	0.01	1% quarterly
Backing capacity	B_{ss}^{short}	7.40	Baseline regime (50 bp fee)
Home productivity	$A_{H,ss}$	1.0	Normalization
Foreign productivity	$A_{F,ss}$	1.0	Normalization
CBDC extension only			
Home USD invoicing	ω_H^{USD}	0.95	Home-currency pricing (total absorption)
Foreign USD invoicing	ω_F^{USD}	0.10	Dominant currency × import share
Network effect	η	1.0	Moderate lock-in

Notes: Parameters except CBDC extension block are common to both models. χ (labor disutility) is endogenous, targeting $N_H = N_F = 0.33$. $\bar{B}$ equals steady-state NFA.

Trade structure: The Armington elasticity $\nu = 1.5$ is within the range of empirical estimates for aggregate trade elasticities. Import shares $\omega_H = \omega_F = 0.25$ imply that each country sources 25% of its final absorption from abroad.

Stablecoin mechanism. The velocity parameter $\theta = 0.5$ governs how much settlement capacity each unit of Treasury backing provides. The fee-function parameters $\kappa = 0.002$ and $p = 2$ are calibrated jointly with θ to produce a steady-state fee around 50 basis points (annualised) under the baseline regime (Section 6.1, $B_{ss}^{short} \approx 7.40$). This target is broadly consistent with the spreads observed in major USD stablecoin payment processing, and the implied utilisation ($u_{ss} \approx 0.54$) sits well below the threshold where fees explode. The convexity parameter $p = 2$ implies quadratic congestion, capturing the realistic feature that marginal settlement costs rise sharply as backing capacity is exhausted.

External adjustment: The risk premium parameter $\phi_B = 0.05$ ensures that net foreign asset positions are stationary: deviations from $\bar{B}$ generate risk premium adjustments that stabilize the current account. The target $\bar{B}$ itself is determined endogenously in steady state to satisfy the NFA accumulation equation given trade flows and fee transfers.

Shock processes: Productivity shocks have persistence $\rho_A = 0.95$, standard in RBC models. Monetary policy shocks (interest rates) have $\rho_i = 0.90$, reflecting central bank gradualism. The backing capacity shock $\rho_{B^{short}} = 0.90$ captures persistence in US Treasury bill supply, which varies with fiscal policy and debt management operations.

CBDC extension parameters: The mandatory USD invoicing shares $\omega_H^{USD} = 0.95$ and $\omega_F^{USD} = 0.10$ are derived from the empirical dollar-invoicing shares in international trade documented by Gopinath and Stein (2021) (approximately 85% for US trade and 40–50% for the rest of the world), rescaled to the total-absorption scope of the model (Section 4.1.1). Foreign domestic transactions are predominantly invoiced in the Foreign currency unit; only the cross-border component of Foreign absorption, equal to the import share $\omega_F = 0.25$, is plausibly USD-invoiced at the Gopinath–Stein rate, giving $\omega_F^{USD} \approx 0.25 \times 0.45 \approx 0.10$. For Home, home-currency pricing yields near-universal USD invoicing across both domestic and import flows, giving $\omega_H^{USD} \approx 0.95$. These shares set the *mandatory* floor for USD stablecoin usage; the residual $(1 - \omega_j^{USD})$ is discretionary and can be diverted to Foreign CBDC. The network effect parameter $\eta = 1.0$ governs the strength of coordination externalities in settlement choice: higher values amplify tipping dynamics toward the dominant settlement system. We set $\eta = 1.0$ to capture moderate but non-negligible network effects, ensuring the model exhibits realistic co-existence of both settlement systems without extreme lock-in.

The two model specifications (baseline with stablecoins only, and CBDC extension) have been defined in Sections 3 and 4 respectively. Computational details on the Dynare implementation and the structural shock processes used to generate the impulse responses are reported in Appendix A.

6 Quantitative results

This section develops the quantitative implications of the safe-asset capacity channel. We first compare steady states across three Treasury-scarcity regimes — calibrated by the equilibrium settlement fee they generate — and document the large permanent welfare effects that follow from moving across regimes (Section 6.1). We then turn to the transmission of shocks around the baseline steady state, where the convex fee schedule is comparatively flat and the spillover channel is dormant (Section 6.2). A brief discussion of the issuance mechanism of the Foreign CBDC closes the section (Section 6.3).

The central finding is that the cross-steady-state comparison — *not* the IRF amplification around a given steady state — is where the main welfare action lies. Under the extreme regime ($fee^{ss} = 300$ bp annualised), the cross-border fee transfer from Foreign to Home reaches almost 3% of Foreign annual GDP, with a Foreign welfare-equivalent consumption loss of roughly 1.3%. A Foreign CBDC reduces this loss by approximately 80%.

6.1 Steady-state landscape across Treasury scarcity regimes

We organise the analysis around three Treasury-scarcity scenarios, each indexed by the annualised stablecoin settlement fee that emerges endogenously in the baseline (no-CBDC) steady state. Anchoring scenarios to the equilibrium fee — rather than to the underlying short-Treasury supply parameter B^{short} — makes the analysis empirically transparent: the fee is the observable manifestation of Treasury scarcity, and serves as a sufficient statistic for the intensity of the safe-asset capacity channel.

Scenario calibration. The three scenarios are:

- **Baseline**: annualised fee of 50 basis points, broadly consistent with current empirical observations on USD stablecoin settlement spreads. Bisection of the steady-state system pins the implied Treasury backing at $B^{short} \approx 7.40$.
- **Stressed**: 150 basis points — a moderate deterioration of Treasury collateral availability, plausibly produced by rapid stablecoin expansion or a partial reduction in short-dated Treasury supply. Implied $B^{short} \approx 5.51$.
- **Extreme**: 300 basis points — a severe scenario where Treasury scarcity binds substantially. Implied $B^{short} \approx 4.84$.

Table 3 reports steady-state outcomes for each scenario, both in the baseline model (USD stablecoins only) and in the CBDC extension where the Foreign central bank issues a costless alternative settlement instrument at the same B^{short}.

Permanent welfare effects. The transition from the baseline to the stressed or extreme regime generates large *permanent* welfare effects, sharply asymmetric across the two countries. Foreign households suffer a permanent reduction in consumption: $\Delta C_F = -0.35\%$ in the stressed scenario and -0.87% in the extreme one. Foreign labour supply N_F *rises* (+0.17% and +0.44%) through the income effect on labour supply: lower consumption raises the marginal utility of consumption, increasing the household's willingness to work. The combination of lower consumption together with higher work effort amplifies the welfare cost: the consumption-equivalent welfare loss for Foreign is -0.51% at the stressed calibration and -1.27% at the extreme calibration.

Table 3: Steady-state outcomes across Treasury scarcity regimes

	Baseline	**Stressed**	**Extreme**
Calibration			
Target annualised SS fee (bp)	50	150	300
Implied B^{short}	7.40	5.51	4.84
u_{ss} (no CBDC)	0.54	0.72	0.82
Steady-state stablecoin fee (annualised bp)			
No CBDC	50.0	150.0	300.0
With CBDC	30.8	74.3	114.7
CBDC reduction (%)	38.4	50.5	61.8
Foreign permanent effects (% deviation vs Baseline)			
ΔC_F, no CBDC	0	-0.35	-0.87
ΔC_F, with CBDC	0	-0.10	-0.19
ΔN_F, no CBDC	0	$+0.17$	$+0.44$
ΔK_F, no CBDC	0	-0.35	-0.87
ΔY_F, no CBDC	0	≈ 0	≈ 0
Welfare-eq. C loss (%), no CBDC	0	-0.51	-1.27
Welfare-eq. C loss (%), with CBDC	0	-0.15	-0.27
CBDC welfare insurance (% reduction)	—	71	79
Home permanent effects (% deviation vs Baseline)			
ΔC_H, no CBDC	0	$+0.21$	$+0.52$
ΔY_H, no CBDC	0	-0.22	-0.53
Welfare-eq. C gain (%), no CBDC	0	$+0.42$	$+1.08$
Digital exorbitant privilege: FEE^F (annualised, % of Y_F^{ss})			
No CBDC	0.50	1.49	2.94
With CBDC	0.22	0.51	0.76

Notes. Scenarios are indexed by the annualised SS fee in the no-CBDC model. Implied B^{short} values are obtained by bisection. The CBDC extension column uses the same B^{short} to isolate the role of CBDC as an alternative settlement instrument. Welfare-equivalent consumption variation is computed as $dCE_j = dC_j - W_j^* \, dN_j$, where $W_F^* = (1-\alpha)Y_F/[N_F(1+fee_{ss})]$ is the after-fee real wage in consumption units; the corresponding Home expression has no fee wedge. ΔY_F is reported as "≈ 0" since the Cobb-Douglas combination of falling K_F and rising N_F nets to less than 0.01% in absolute value.

Home households experience the mirror image: C_H rises by +0.21% and +0.52%, N_H falls (Home works less because of the implicit income transfer from Foreign), and the consumption-equivalent welfare gain is +0.42% and +1.08%. The gain is somewhat smaller in magnitude than Foreign's loss, reflecting a modest net global welfare loss — but the cross-country redistribution dominates.

The Foreign output puzzle. A perhaps surprising feature of Table 3 is that Foreign output Y_F is essentially unchanged across scenarios. The mechanism is the following. A higher fee distorts the Foreign capital Euler through the wedge $(1 + fee)$, so the Foreign capital stock K_F falls (−0.35% stressed, −0.87% extreme). But Foreign labour supply rises through the income effect. With Cobb-Douglas technology, $\Delta \log Y_F = \alpha \, \Delta \log K_F + (1-\alpha) \, \Delta \log N_F$. With $\alpha = 0.33$ and the two movements approximately offsetting in our calibration, $\Delta Y_F \approx 0$. The welfare cost of Treasury scarcity therefore manifests in the *composition* of the Foreign equilibrium — less capital, more labour effort, lower consumption — rather than in its overall scale.

The size of the digital exorbitant privilege. The cross-border fee transfer FEE^F from Foreign to Home grows substantially with Treasury scarcity. As a fraction of Foreign annualised output, it equals 0.50% at the baseline, 1.49% in the stressed scenario, and 2.94% in the extreme one. At the extreme regime, this implies that nearly 3% of Foreign GDP flows permanently to Home as compensation for using the USD-denominated settlement infrastructure — a structural transfer that, absent an alternative settlement system, would not be politically sustainable.

The macroprudential insurance value of CBDC. Comparing the baseline model with the CBDC extension at the same B^{short} reveals the substantial insurance value of a Foreign CBDC. The Foreign central bank's costless alternative settlement instrument absorbs marginal demand and prevents the stablecoin fee from rising to its no-CBDC level: at the stressed calibration the fee is reduced from 150 to 74 bp (50% reduction); at the extreme calibration from 300 to 115 bp (62% reduction). The corresponding reductions in the Foreign welfare-equivalent loss are even larger — 71% at stressed and 79% at extreme — because the wedge enters several margins of the Foreign household problem and the nonlinear fee schedule is steepest where the CBDC absorbs the most demand. The annual fee transfer is similarly compressed: from 1.49% to 0.51% of Foreign output at stressed; from 2.94% to 0.76% at extreme. The insurance value of CBDC is strongly convex in the severity of Treasury scarcity — modest under benign conditions (roughly a third of the fee absorbed) and substantial under stress (close to two-thirds at extreme), exactly the property one would want of a macroprudential instrument targeting settlement infrastructure congestion.

Roadmap for the remainder of this section. Section 6.2 characterises the transmission of shocks around the baseline steady state, where the convex fee schedule is relatively flat and the spillover channel is dormant. Section 6.3 discusses the issuance mechanism of the Foreign CBDC.

6.2 Transmission of shocks around the baseline steady state

We now characterise the impulse responses to the five structural shocks of the model around the baseline steady state, where the convex fee schedule is relatively flat and the safe asset capacity channel is dormant. The aim is to identify the channels through which shocks propagate to the settlement layer and through which the settlement layer feeds back into real activity, leaving the magnification of these channels at more congested steady states to Section 6.1. A notable feature of our specification, relative to symmetric-normalisation formulations, is the endogenous Foreign producer price $P_{FF,t}$: when Foreign productivity or exchange-rate conditions change, $P_{FF,t}$ adjusts, compressing or expanding Foreign invoice values $INV_F^{USD} = P_{F,t}X_{F,t}/E_t$ directly. This *cost-deflation channel* attenuates the settlement-demand response to productivity shocks.

Appendix B reports the impulse responses to one-standard-deviation shocks to, respectively, the US interest rate i^{US}, the Foreign interest rate i^F, the US short-Treasury supply B^{short}, Home productivity A_H, and Foreign productivity A_F (Figures 2–6). Each figure overlays the baseline (no-CBDC) and CBDC-extension responses. The shocks are calibrated to standard deviations of 0.002 for nominal rates and 0.01 (1%) for the three real-side shocks. Table 4 summarises the peak fee response for each shock.

Three robust messages emerge from these IRFs. First, fee deviations under the baseline

Table 4: Peak settlement-fee response by shock, baseline regime

Shock	No CBDC (bp ann)	With CBDC (bp ann)	IRF figure
US interest rate (+1 s.d.)	−0.34	−0.09	2
Foreign interest rate (+1 s.d.)	+0.34	+0.09	3
Backing capacity (+1%)	−1.18	−0.67	4
US productivity (+1%)	+1.02	+0.62	5
Foreign productivity (+1%)	+0.33	+0.15	6

Notes. IRF figures are collected in Appendix B, where shock-by-shock commentary on each panel is also reported.

calibration are uniformly small (at most ±1.2 basis points annualised), confirming that the safe-asset capacity channel is dormant when utilisation is in the flat region of the fee schedule. Real-side output, consumption, and investment responses are driven overwhelmingly by standard RBC channels — terms of trade, real wages, intertemporal substitution — and the baseline and CBDC impulse responses are virtually indistinguishable along these dimensions. Second, the asymmetric incidence of the settlement fee across the two countries (a circular transfer for Home, a genuine resource wedge for Foreign) generates a small but consistent Foreign-response premium: Foreign output reacts more strongly than Home to direct shocks on the settlement layer, such as the B^{short} shock, because the relaxation of the Foreign-side wedge stimulates Foreign consumption, investment, and labour supply directly while Home gains only through standard general-equilibrium spillovers. Third, the cost-deflation channel materially attenuates the fee response to productivity shocks — the Foreign productivity fee response is roughly one-third of the Home one, despite the two shocks having equal magnitude, because the endogenous fall in $P_{FF,t}$ compresses Foreign invoice values and partly offsets the inflationary push from dollar depreciation.

6.3 The issuance mechanism of CBDC

A natural question is whether our results depend on *how* the Foreign central bank issues the CBDC. In a closed-economy DSGE with a retail CBDC backed by domestic sovereign debt, Barrdear and Kumhof (2022) show that the central bank's outright bond purchases compress the supply of safe assets available to private investors, generating real effects through lower equilibrium real rates and reduced distortionary taxation. Could a similar channel undo our results by tightening the safe-asset constraint that drives the stablecoin fee schedule?

The answer is no, for two reasons. First, the most advanced CBDC project at the time of writing — the digital euro — is not designed as an asset-purchase programme. When a household converts a bank deposit into digital euro, the converting bank's reserve account at the central bank is debited and the household's digital-euro balance is credited, leaving the central bank's asset side unchanged and only shifting the composition of its liabilities from reserves to public-money holdings. Any aggregate reserve drain that ensues is typically accommodated through standard collateralised refinancing operations (i.e., lending to banks), not through outright purchases of sovereign debt; in addition, the announced holding caps bound the scale of conversion precisely to discourage balance-sheet expansion.

Second, and more fundamentally for our framework, the safe asset that backs the stablecoin in our model is the stock of *Home* (US) Treasuries, while the Foreign CBDC — regardless of its issuance mechanism — is a liability of the Foreign central bank, backed ultimately by Foreign-currency assets. Even in the hypothetical scenario in which the Foreign central bank financed CBDC issuance through outright purchases of Foreign sovereign debt, those purchases would compress the supply of Bunds, OATs or BTPs, not US Treasuries. The safe-asset capacity channel that drives our results is therefore robust to the issuance design of the Foreign CBDC.

The Barrdear-Kumhof channel *would* become operative if the analysis were extended to a hypothetical *Home* (USD) CBDC issued against US Treasury purchases. In that case the Federal Reserve would simultaneously (i) provide an alternative settlement instrument that relieves stablecoin demand and (ii) compress the Treasury collateral pool that backs private stablecoins, generating an internal trade-off that is absent from the Foreign-CBDC exercise studied here. We leave this domestic-CBDC counterpart to future work.

7 Robustness

Section 6.1 already documented the cross-steady-state welfare implications of varying the Treasury backing capacity B^{short}. This section assesses the sensitivity of the model to the remaining structural parameters, anchoring all comparisons to the baseline regime of Table 3 (annualised SS fee ≈ 50 bp). Table 5 reports the implied steady-state utilisation and fee for both the baseline and the CBDC extension across alternative values of these parameters; the narrative that follows discusses the corresponding permanent welfare implications by analogy with Section 6.1.

7.1 Collateral efficiency (θ)

The parameter θ governs how efficiently safe assets are intermediated into settlement capacity: the effective backing is $\theta \cdot B^{\text{short}}$. Because θ and B^{short} enter the utilisation formula symmetrically, a change in θ moves the economy across regimes analogously to a change in B^{short}. Reducing θ from 0.5 to 0.3 raises steady-state utilisation enough to push the SS fee well above the extreme regime of Table 3 (300 bp), with correspondingly larger permanent welfare losses; raising θ to 0.8 produces a fee comfortably below the baseline regime, with negligible welfare effects. Improvements in collateral intermediation efficiency — better repo market infrastructure, reduced regulatory haircuts on Treasury collateral, or more efficient intraday liquidity management by stablecoin issuers — are therefore equivalent to expanding the effective supply of backing assets and can move the economy across welfare regimes without any change in Treasury issuance.

7.2 Fee schedule parameters (κ, p)

The fee schedule $\mathit{fee} = \kappa \cdot u^p/(1-u)$ has two degrees of freedom: the scale parameter κ and the convexity exponent p.

Fee scale (κ). Raising κ increases the fee level approximately proportionally without affecting utilisation. Moving κ from the baseline value 0.002 to 0.01 scales the SS fee linearly from the baseline regime to roughly 200 bp (between stressed and extreme); halving κ to 0.001 halves the

Table 5: Sensitivity of steady-state outcomes to key parameters

	Baseline (no-CBDC) model		CBDC extension			
	u	Fee (bp)	u	Fee (bp)	$share_F^{USD}$	CBDC %
Collateral efficiency θ						
$\theta = 0.3$	0.84	342	0.68	118	0.67	17.5
$\theta = 0.5$ (baseline)	0.50	40	0.42	25	0.72	14.8
$\theta = 0.8$	0.31	11	0.27	8	0.73	14.3
Fee scale κ						
$\kappa = 0.001$	0.50	20	0.43	13	0.73	14.5
$\kappa = 0.002$ (baseline)	0.50	40	0.42	25	0.72	14.8
$\kappa = 0.01$	0.50	200	0.41	115	0.67	17.4
Fee convexity p						
$p = 1$	0.50	80	0.42	58	0.70	15.7
$p = 2$ (baseline)	0.50	40	0.42	25	0.72	14.8
$p = 4$	0.50	10	0.43	5	0.73	14.3
Foreign invoicing share ω_F^{USD} (CBDC extension only)						
$\omega_F^{USD} = 0.05$	—	—	0.42	24	0.70	15.7
$\omega_F^{USD} = 0.10$ (baseline)	—	—	0.42	25	0.72	14.8
$\omega_F^{USD} = 0.20$	—	—	0.43	26	0.75	13.1
$\omega_F^{USD} = 0.45$	—	—	0.45	30	0.83	9.1
Network effect η (CBDC extension only)						
$\eta = 0$	—	—	0.38	18	0.54	24.3
$\eta = 1$ (baseline)	—	—	0.42	25	0.72	14.8
$\eta = 2$	—	—	0.46	32	0.87	6.8
$\eta = 5$	—	—	0.50	39	0.99	0.4

Notes. Each row varies one parameter from the baseline calibration $\{B^{short} = 8, \theta = 0.5, \kappa = 0.002, p = 2, \omega_F^{USD} = 0.10, \eta = 1\}$ (indicated by "baseline" labels), holding all others fixed. This calibration yields an SS fee of ≈ 40 bp, slightly below the 50 bp baseline regime of Table 3 (obtained by bisecting B^{short} to 7.40); the qualitative robustness rankings are unchanged. Fees are annualised basis points. Baseline model columns are blank for parameters that affect only the CBDC extension (ω_F^{USD}, η). The robustness exercise on B^{short} is reported in Section 6.1 and is not repeated here.

fee. The permanent welfare implications track the regime mapping of Table 3. For the CBDC model, a higher κ raises the fee differential between USD and CBDC settlement, inducing more substitution and a larger insurance value along all regimes.

Fee convexity (p). The exponent p controls the shape of the fee schedule. Lowering p to 1 (linear numerator) approximately doubles the SS fee relative to the baseline regime and modestly amplifies IRF fee responses; raising p to 4 compresses the fee to well below the baseline regime, with correspondingly muted IRF responses. Higher p concentrates fee sensitivity at high utilisation, making the model more threshold-like. Our choice of $p = 2$ represents a middle ground consistent with the empirical observation that repo spreads begin to widen gradually before capacity constraints bind severely.

7.3 Invoicing shares ($\omega_H^{\mathbf{USD}}$, $\omega_F^{\mathbf{USD}}$)

The mandatory invoicing parameters determine the floor of USD settlement that cannot be diverted to CBDC regardless of fee differentials. They are specific to the CBDC extension and govern the maximum scope for substitution.

Under our baseline calibration $\omega_F^{\mathrm{USD}} = 0.10$ (justified in Section 4.1.1), 90% of Foreign absorption is in principle eligible for migration to a foreign CBDC. Network externalities ($\eta = 1.0$) nevertheless sustain $share_F^{\mathrm{USD}} \approx 0.72$ in equilibrium. Raising the mandatory floor chips away at the discretionary margin but does not by itself dislodge the network-effect coordination equilibrium: even at $\omega_F^{\mathrm{USD}} = 0.45$ — the value a naive reading of the cross-border invoicing data would suggest — the equilibrium USD share rises only to 0.83.

The cross-steady-state welfare implications are also modest. Because ω_F^{USD} affects only the CBDC extension, raising it from 0.10 to 0.45 increases the SS fee from 25 to 30 basis points (at the baseline B^{short}), a small change in regime terms; the corresponding increase in the Foreign welfare-equivalent loss is of order 0.03 percentage points. The binding constraint on CBDC effectiveness is the network channel rather than the invoicing channel, and policy levers focused on the latter (e.g. promotion of euro or renminbi invoicing) deliver only second-order gains.

7.4 Network effect strength (η)

The network effect parameter η governs the strength of the coordination externality that favours incumbent USD stablecoins. Its effect on the CBDC equilibrium is dramatic.

At $\eta = 0$ (no network effects), the CBDC share rises to 24% of total settlement and the equilibrium SS fee falls to 18 basis points — below the baseline regime, with negligible permanent Foreign welfare loss. At $\eta = 2$, network lock-in dominates: the CBDC share shrinks to 7%, the fee rises to 32 basis points, and the USD settlement shares climb to $share_H^{USD} = 0.98$ and $share_F^{USD} = 0.87$. By $\eta = 5$, CBDC is virtually irrelevant: its share falls to about 0.4% and the fee converges to within approximately 0.5 basis points of the no-CBDC value. The practical implication is that at fixed B^{short}, the insurance value of CBDC — and thus the welfare gain from its availability — shrinks rapidly with network strength.

The economic intuition is that network effects create a self-reinforcing equilibrium: more USD stablecoin usage raises the network benefit, which attracts further adoption, which raises the benefit further. A strong network effect ($\eta \geq 3$) effectively eliminates CBDC as a competitive force even when CBDC is costless and unconstrained. This result speaks directly to the policy debate around CBDC adoption: absent deliberate policy interventions that strengthen CBDC network effects — such as mandating CBDC acceptance by financial institutions, subsidising CBDC adoption, or requiring government payments to use CBDC — private stablecoins retain incumbency advantages that a costless public alternative cannot easily dislodge.

More generally, the value of CBDC as a tail-risk insurance mechanism depends not only on backing capacity conditions but also on the degree to which network effects have already locked in USD dominance: in the stressed and extreme regimes of Table 3, the welfare-equivalent insurance gains from CBDC (71% and 79% reductions) hinge on the moderate network calibration $\eta = 1$ and would shrink materially under stronger lock-in.

8 Conclusions

This paper develops a two-country RBC model in which USD stablecoins, backed by short-term US Treasuries, serve as the dominant global settlement technology. The convex relation between stablecoin utilisation and the equilibrium settlement fee — driven by safe-asset backing capacity — generates a new international spillover channel that operates on two reinforcing margins: a permanent welfare margin across alternative steady states, and a dynamic margin around any given steady state.

Permanent welfare effects and the digital exorbitant privilege. The cross-steady-state comparison is where the headline action lies. Moving from a baseline regime (50 bp annualised fee) to stressed and extreme regimes (150 and 300 bp) reduces Foreign consumption permanently by 0.35% and 0.87%, with consumption-equivalent welfare losses of 0.51% and 1.27%. The cross-border fee transfer from Foreign to Home — which we interpret as a *digital exorbitant privilege* mirroring the classical version through ownership of settlement infrastructure rather than yield differentials on the external balance sheet — rises from 0.5% to 2.9% of Foreign annual GDP across the three regimes. The output puzzle is informative in its own right: Foreign output is essentially flat across regimes, while consumption falls and labour supply rises through the income effect. The welfare cost of safe-asset scarcity therefore manifests in the composition of the Foreign equilibrium rather than in its scale.

The dynamic margin: a threshold phenomenon. Around any given steady state, the safe asset capacity channel is threshold-dependent. At the baseline regime ($u_{ss} \approx 0.54$) the convex fee schedule is nearly flat: shocks generate fee deviations of at most ±1.2 basis points (annualised), and real responses are dominated by standard RBC channels. As one moves to the stressed and extreme regimes, the slope of the fee schedule rises roughly threefold and eightfold respectively, activating three reinforcing transmission channels — wealth transfer, intertemporal distortion, and congestion feedback. Dollar exchange-rate movements interact asymmetrically with this layer: appreciation shrinks the USD value of Foreign invoices and relieves settlement pressure, while depreciation amplifies it.

CBDC as nonlinear insurance. A costless Foreign CBDC provides state-dependent insurance whose value is strongly convex in safe-asset scarcity. SS fee reduction rises from 38% at the baseline regime to 51% (stressed) and 62% (extreme), while the corresponding Foreign welfare-equivalent loss reduction reaches 71% and 79%. CBDC is largely irrelevant under normal conditions but becomes decisive under stress. The binding constraint on CBDC effectiveness is the strength of network externalities favouring the incumbent settlement system, not the mandatory invoicing floor.

Policy implications. Three implications follow. First, the relevant financial-stability metric for stablecoin-based international settlement is not stablecoin market capitalisation in isolation but the ratio of issuers' Treasury-instrument holdings to the outstanding stock of short-dated US Treasury securities — the structural parameter that governs both the permanent welfare cost and the dynamic IRF amplification. Second, the first-best response to stablecoin growth is to ensure that short-term Treasury supply expands broadly in line with adoption, keeping utilisation in the flat regime where spillovers are negligible. Third, CBDC should be evaluated as a tail-risk insurance mechanism: its welfare value is small under normal conditions but large under stress, and its effectiveness depends primarily on the network-coordination equilibrium between incumbent USD stablecoins and the alternative settlement instrument rather than on mandatory invoicing constraints. Policies that strengthen the network value of a competing public settlement system — through mandatory acceptance, government payment requirements, or interoperability standards — can shift this coordination equilibrium in ways that incremental changes to invoicing conventions alone cannot.

References

Ahmed, R., & Aldasoro, I. (2025). Stablecoins and safe asset prices. *BIS Working Papers*, No. 1270.

Auer, R., Haene, P., & Holden, H. (2021). Multi-CBDC arrangements and the future of cross-border payments. *BIS Papers*, No. 115.

Azzimonti, M., & Quadrini, V. (2025). Digital Economy, Stablecoins, and the Global Financial System. NBER Working Paper 34066.

Barrdear, J., & Kumhof, M. (2022). The macroeconomics of central bank digital currencies. *Journal of Economic Dynamics and Control*, 142, 104148.

Benigno, P., Schilling, L. M., & Uhlig, H. (2022). Cryptocurrencies, currency competition, and the impossible trinity. *Journal of International Economics*, 136, 103601.

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

Caballero, R. J., & Farhi, E. (2018). The safety trap. *Review of Economic Studies*, 85(1), 223–274.

Caballero, R. J., Farhi, E., & Gourinchas, P.-O. (2008). An equilibrium model of "global imbalances" and low interest rates. *American Economic Review*, 98(1), 358–393.

Catalini, C., & de Gortari, A. (2021). On the economic design of stablecoins. MIT Cryptoeconomics Lab Working Paper.

Cong, L. W., Li, Y., & Wang, N. (2022). Token-based platform finance. *Journal of Financial Economics*, 144(3), 972–991.

Cova, P., Notarpietro, A., Pagano, P., & Pisani, M. (2026). Monetary policy in the open economy with digital currencies. *International Finance*, 29(1), 79–91.

Davoodalhosseini, M. (2022). Central bank digital currency and monetary policy. *Journal of Economic Dynamics and Control*, 142, 104150.

Fernández-Villaverde, J., Sanches, D., Schilling, L., & Uhlig, H. (2021). Central bank digital currency: Central banking for all? *Review of Economic Dynamics*, 41, 225–242.

Ferrari Minesso, M., & Siena, D. (2026). Stablecoins and the global safe asset market. Working Paper.

Gopinath, G., Boz, E., Casas, C., Díez, F. J., Gourinchas, P.-O., & Plagborg-Møller, M. (2020). Dominant currency paradigm. *American Economic Review*, 110(3), 677–719.

Gopinath, G., & Stein, J. C. (2021). Banking, trade, and the making of a dominant currency. *Quarterly Journal of Economics*, 136(2), 783–830.

Gorton, G., & Zhang, J. (2023). Taming wildcat stablecoins. *University of Chicago Law Review*, 90(3), 909–971.

Gourinchas, P.-O., & Rey, H. (2007). From world banker to world venture capitalist: U.S. external adjustment and the exorbitant privilege. In R. H. Clarida (Ed.), *G7 Current Account Imbalances: Sustainability and Adjustment* (pp. 11–66). University of Chicago Press.

Gourinchas, P.-O., Rey, H., & Govillot, N. (2019). Exorbitant privilege and exorbitant duty. IMES Discussion Paper Series 2019-E-1, Bank of Japan.

He, Z., Krishnamurthy, A., & Milbradt, K. (2019). A model of safe asset determination. *American Economic Review*, 109(4), 1230–1262.

Jiang, Z., Krishnamurthy, A., & Lustig, H. (2021). Foreign safe asset demand and the dollar exchange rate. *Journal of Finance*, 76(3), 1049–1089.

Krishnamurthy, A., & Vissing-Jorgensen, A. (2012). The aggregate demand for Treasury debt. *Journal of Political Economy*, 120(2), 233–267.

Le, A., et al. (2023). Digital dollarization in small open economies. Working Paper.

Lyons, R. K., & Viswanath-Natraj, G. (2023). What keeps stablecoins stable? *Journal of International Money and Finance*, 131, 102777.

Maggiori, M., Neiman, B., & Schreger, J. (2020). International currencies and capital allocation. *Journal of Political Economy*, 128(6), 2019–2066.

Makarov, I., & Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2), 293–319.

Makarov, I., & Schoar, A. (2022). Cryptocurrencies and decentralized finance (DeFi). *Brookings Papers on Economic Activity*, Spring 2022, 141–196.

Routledge, B., & Zetlin-Jones, A. (2022). Currency stability using blockchain technology. *Journal of Economic Dynamics and Control*, 142, 104155.

Schmitt-Grohé, S., & Uribe, M. (2003). Closing small open economy models. *Journal of International Economics*, 61(1), 163–185.

A Computational details

This appendix collects the technical material on the Dynare implementation and on the structural shock processes used in the impulse-response analysis.

Solver and model dimensions. Both models are solved using Dynare 7.0. The baseline model contains 44 endogenous variables and 5 exogenous shocks; the CBDC extension contains 54 endogenous variables with the same shock structure.[5] Steady states are computed using a combination of analytical conditions (Euler equations, factor demands) and numerical fixed-point iteration for the fee and settlement shares. In the CBDC model, we iterate on

$$\begin{cases} S_t = (1 + \mathit{fee}_t) \cdot \big(\mathit{share}^{USD}_{H,t} \cdot \mathit{INV}^{USD}_{H,t} + \mathit{share}^{USD}_{F,t} \cdot \mathit{INV}^{USD}_{F,t}\big) \\ \mathit{fee}_t = \kappa \, \dfrac{u_t^p}{1 - u_t + \bar{\varepsilon}}, \quad u_t = \dfrac{S_t}{\theta \cdot B^{short}} \\ \mathit{share}^{USD}_{j,t} = \omega_j^{USD} + \dfrac{1 - \omega_j^{USD}}{1 + \exp\Big[-\eta \cdot \dfrac{S_t}{S_t + CBDC_t^{USD} + \bar{\varepsilon}} \; + \; \psi \cdot \mathit{fee}_t\Big]} \end{cases}$$

where $\bar{\varepsilon} > 0$ is a small numerical safeguard and $\psi = 100$ rescales the fee into units comparable to the network term. The Blanchard–Kahn conditions are satisfied in both specifications. The baseline model's slightly asymmetric steady state, arising from $K_{F,ss} < K_{H,ss}$ through the fee distortion in the Foreign Euler, is computed numerically using `steady(solve_algo=4)` starting from the standard RBC symmetric approximation; the correction is of order $\mathit{fee}_{ss} \approx 0.1\%$.

Shock processes. The five structural shocks have first-order persistence parameters $\rho_{A_H} = \rho_{A_F} = 0.95$, $\rho_{i^{US}} = \rho_{i^F} = 0.90$ and $\rho_{B^{short}} = 0.90$. Their standard deviations are calibrated as follows: $\sigma_{i^{US}} = \sigma_{i^F} = 0.002$ (20 annualised basis points), and $\sigma_{B^{short}} = \sigma_{A_H} = \sigma_{A_F} = 0.01$ (a 1% shock). Impulse responses are computed at first order over a 40-quarter horizon and overlaid for the baseline and CBDC specifications to isolate the stabilisation effects of CBDC availability.

B Detailed impulse-response analysis at the baseline regime

This appendix reports the full shock-by-shock commentary on the impulse-response functions previewed in Section 6.2. All numerical magnitudes are computed at the baseline regime ($u_{ss} \approx 0.54$, SS fee 50 basis points annualised); the cross-steady-state magnification of these responses at the stressed and extreme regimes follows from the fee-schedule slopes reported in Section 6.1.

---

[5] Many of these are static definitions or reporting identities (e.g. Armington demands, CES price indices, national-accounts aggregates). The "core" dynamic variables — those that appear with leads or lags — number 17 in both models: consumption, investment, capital, labour, output, and TFP for each country, the exchange rate, the two policy rates, the bond supply, and net foreign assets.

B.1 US interest rate shock

Figure 2 displays the impulse responses to a one standard deviation positive shock to the US interest rate ($e_{i^{US}}$). This shock creates a differential between Home and Foreign returns, triggering an immediate appreciation of the dollar (E rises by approximately 0.55%) via the UIP condition.

The dollar appreciation induces standard expenditure switching: Foreign goods become cheaper for Home consumers while Home goods become more expensive abroad. Home imports rise ($X_{FH} \uparrow$) and exports fall ($X_{HF} \downarrow$), deteriorating the Home trade balance. Despite this, Home output (Y_H) expands by approximately 0.07% (peaking around quarter 4) while Foreign output (Y_F) contracts by the same magnitude. The dominant force is the supply-side impact of the terms of trade on labour decisions: the stronger dollar lowers P_H relative to the Home producer price, raising the real consumption wage $(1-\alpha)Y_H/(N_H P_H)$ and stimulating Home labour supply. For Foreign households, the weaker currency raises P_F, eroding real wages and contracting labour supply and output.

The exchange rate also interacts with the stablecoin settlement layer. Dollar appreciation mechanically reduces the USD value of Foreign invoices ($P_F X_F/E$), relieving pressure on short-term Treasury backing ($u \downarrow$). The equilibrium settlement fee declines by approximately 0.34 basis points (annualised) in the baseline and 0.09 basis points in the CBDC extension — a small but correctly signed response. This "valuation channel" provides a mild stabilising feedback: a strong dollar acts as a liquidity provider for the global settlement system. However, the real effects of this fee movement are negligible: as shown in the output panels, the baseline and CBDC responses are nearly indistinguishable, confirming that fee fluctuations of this magnitude do not materially affect real allocations.

B.2 Foreign interest rate shock

Figure 3 presents the impulse responses to a positive shock to the Foreign interest rate (e_{i^F}). By UIP, the higher return on Foreign assets triggers an appreciation of the Foreign currency, corresponding to a dollar depreciation (E falls by approximately 0.55%).

The transmission mechanism is the mirror image of the US interest rate shock. The weaker dollar raises P_H, depressing the Home real consumption wage and causing Home agents to reduce consumption, investment, and labour supply. Home output contracts by approximately 0.07% (trough around quarter 4). Foreign agents experience the opposite: appreciation raises their real purchasing power, stimulating consumption, investment, and output.

On the settlement side, dollar depreciation inflates the USD value of Foreign invoices ($P_F X_F/E \uparrow$), raising global stablecoin demand. The equilibrium fee rises by approximately 0.34 basis points (annualised) in the baseline and 0.09 basis points in the CBDC extension. As with the US interest rate shock, this fee response is too small to generate first-order real effects: baseline and CBDC output responses are visually indistinguishable.

B.3 Backing capacity shock

Figure 4 displays the impulse responses to a positive 1% shock to short-term Treasury supply ($e_{B^{short}}$). This shock directly expands the collateral pool available for stablecoin backing, leading to an immediate and persistent decline in the equilibrium settlement fee of approximately 1.18 basis points (annualised) in the baseline and 0.67 basis points in the CBDC extension.

The macroeconomic adjustments are of the same sign in both countries but very small, and — consistent with the ownership structure of the stablecoin issuer — Foreign output responds more strongly than Home. In the baseline, Foreign output expands by approximately $1.0 \times 10^{-3}\%$ while Home output expands by only approximately $6 \times 10^{-4}\%$; the CBDC extension produces similarly-signed but smaller responses ($5 \times 10^{-4}\%$ for Foreign and $3 \times 10^{-4}\%$ for Home). These magnitudes reflect the core result: with utilisation at 54% and the fee schedule relatively flat, even a direct shock to backing capacity produces only negligible real effects. The real economy is dominated by standard channels; the settlement layer operates in the background.

The *ranking* of the output responses reflects the asymmetric incidence of the settlement fee. The fee enters as a circular transfer for Home (paid as a consumer, received back as shareholder in the issuer), so it does not distort Home's labour supply or capital accumulation decisions. For Foreign, by contrast, the fee is a genuine resource cost that raises the effective price of Foreign absorption by a factor $(1+fee_t)$ in the Foreign labour FOC and Euler condition. When the fee falls following a Treasury expansion, this effective wedge on Foreign absorption shrinks, stimulating Foreign consumption, investment, and labour supply — Foreign output therefore rises directly and meaningfully. Home output also expands, but only through a general-equilibrium channel: higher Foreign demand for Home goods, combined with a mild dollar depreciation ($E \downarrow$ by approximately $2 \times 10^{-3}\%$), supports Home exports. The backing capacity shock thus raises aggregate efficiency by shrinking a distortionary wedge that applies exclusively to Foreign, rather than merely redistributing welfare across countries.

The utilisation rate panel warrants a separate remark. The utilisation response in Figure 4 shows a trough of approximately 200 basis points, noticeably larger than a naive back-of-the-envelope calculation from the fee response would suggest. This amplification arises from the endogenous $P_{FF,t}$ channel: as the fee falls, Foreign goods prices adjust slightly, compressing Foreign invoice values and further reducing settlement demand in a second-round effect absent from the symmetric-normalisation specification. While the utilisation response is amplified, the real consequences remain negligible because the level of utilisation stays well within the flat region of the fee schedule.

B.4 Home productivity shock

Figure 5 displays the impulse responses to a positive 1% Home productivity shock (e_{A_H}). The shock triggers a robust US expansion: Home output rises to a peak of approximately 1.14%, with Home consumption and investment rising accordingly.

The abundance of Home goods deteriorates the terms of trade, manifesting as a dollar depreciation ($E \downarrow$ by approximately 0.55%). Two forces then simultaneously raise global stablecoin demand. First, higher US real activity directly increases Home invoice values ($INV_H =$

$P_H X_H \uparrow$). Second, the weaker dollar inflates the USD value of Foreign invoices ($INV_F^{USD} = P_F X_F / E \uparrow$). Together these push total utilisation upward, causing the settlement fee to rise by approximately 1.02 basis points (annualised) in the baseline and 0.62 basis points in the CBDC extension — the largest fee increase across all five shocks at the baseline regime. Foreign output rises slightly (approximately +0.02%, peaking around quarter 20).

Under the Home productivity shock, the settlement fee rises. The fee acts as a transfer from Foreign stablecoin users to the Home-based issuer, reducing Foreign disposable income. This negative wealth effect on Foreign households induces them to supply more labour to compensate for the loss in income, which raises Foreign employment and output. The net effect is a small positive Foreign spillover rather than the small negative one which would prevail under a symmetric normalisation. Crucially, in both specifications the Foreign output response is quantitatively negligible at the baseline regime: the settlement fee simply does not generate first-order real effects when the schedule is flat.

B.5 Foreign productivity shock

Figure 6 presents the impulse responses to a positive 1% Foreign productivity shock (e_{A_F}). Foreign output rises to a peak of approximately 1.15%.

The Foreign productivity boom lowers the endogenous Foreign producer price $P_{FF,t}$, making Foreign goods cheaper in FCU terms. With an Armington elasticity of $\nu = 1.5$, Home import demand responds more than proportionally to the fall in the relative price of Foreign goods, generating a significant increase in Home demand for FCU. This competitiveness effect dominates the countervailing income effect (wealthier Foreign households demanding more Home goods), so the FCU appreciates — equivalently, the dollar depreciates ($E \downarrow$ by approximately 0.36% at its peak, around quarter 13).

Home output rises slightly (approximately +0.03%, peaking around quarter 20), driven by the same fee-transfer wealth channel: the higher settlement fee raises Home issuer profits, generating a small positive wealth effect on Home households. Both output responses are quantitatively negligible — the real economy is governed by standard RBC forces.

The settlement fee rises by approximately 0.33 basis points (annualised) in the baseline and 0.15 basis points in the CBDC extension — roughly one-third the response generated by the Home productivity shock of equal size. Two forces push Foreign invoice values $INV_F^{USD} = P_{F,t} X_{F,t} / E_t$ upward: higher Foreign absorption $X_{F,t}$ and dollar depreciation ($E_t \downarrow$) inflating the USD value of Foreign invoices through the denominator. However, the endogenous fall in $P_{FF,t} = \bar{W}_F N_{F,t} / ((1-\alpha) Y_{F,t})$ — as rising Foreign output reduces the labour-to-output ratio — compresses $P_{F,t}$ and partially offsets these inflationary forces through the numerator. The net result is a smaller increase in INV_F^{USD} than under the Home productivity shock, where dollar depreciation and real expansion push invoice values upward with no offsetting compression of Foreign producer prices.

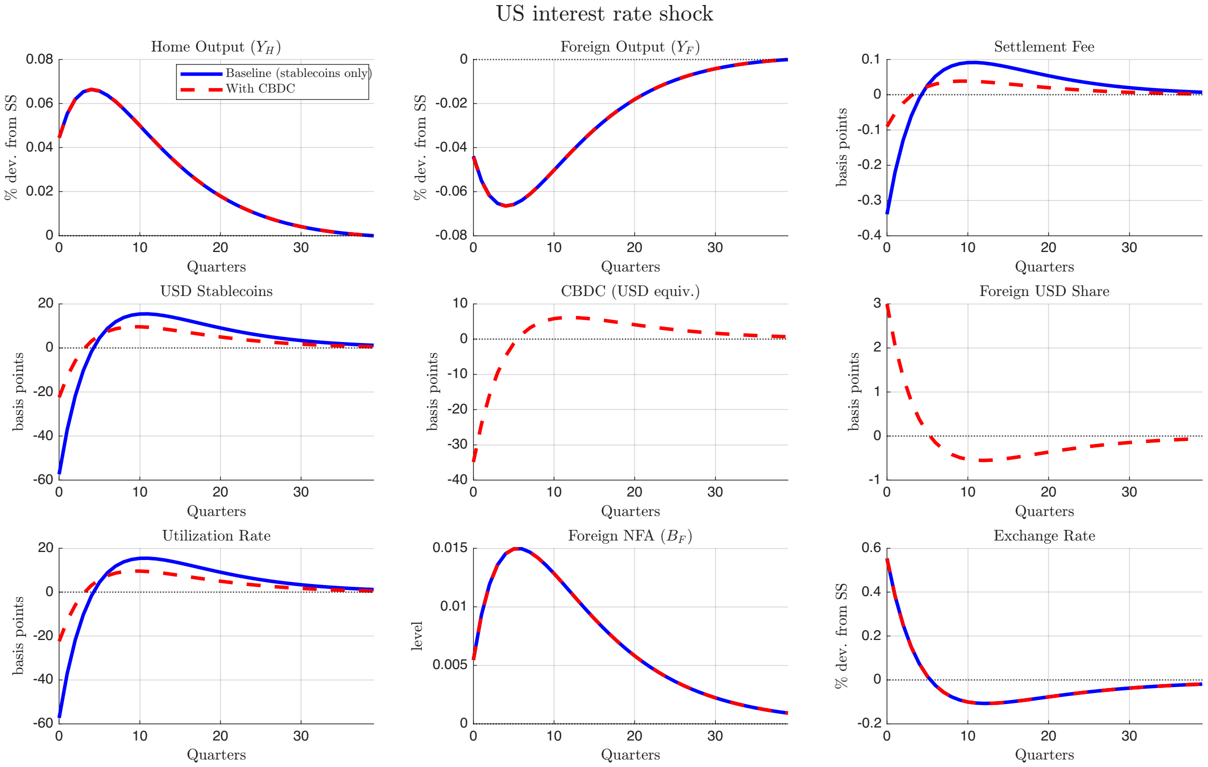

Figure 2: Impulse responses to US interest rate shock.

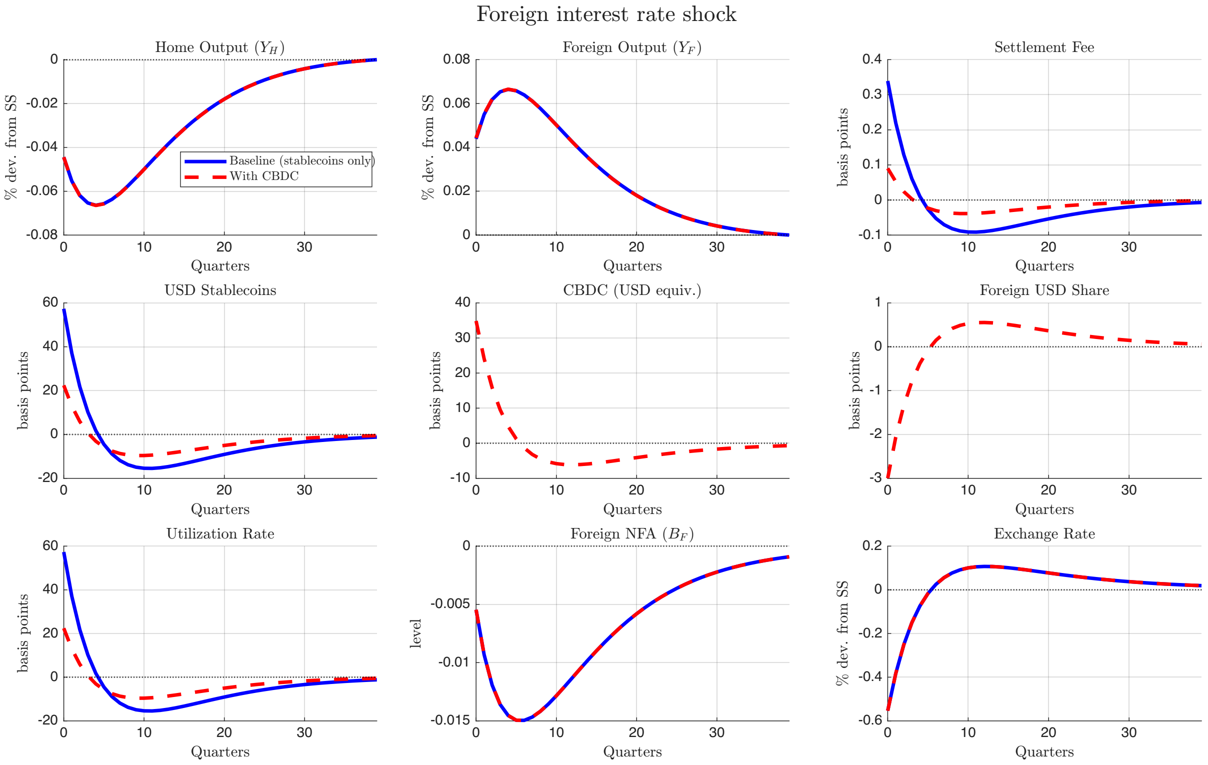

Figure 3: Impulse responses to foreign interest rate shock.

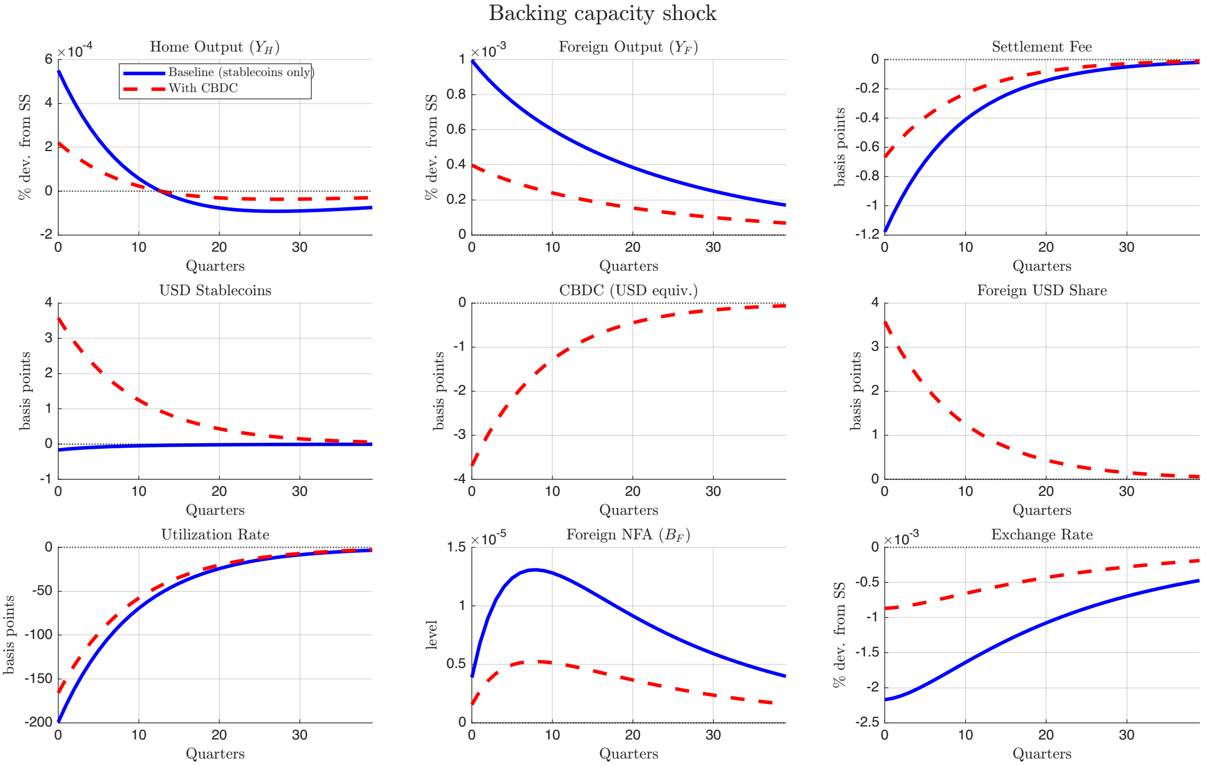

Figure 4: Impulse responses to US backing capacity shock. Output panels are scaled by 10^{-4}; real effects are negligible under the baseline calibration. The large utilisation response (-200 bp) reflects amplification from the endogenous $P_{FF,t}$ channel; see text.

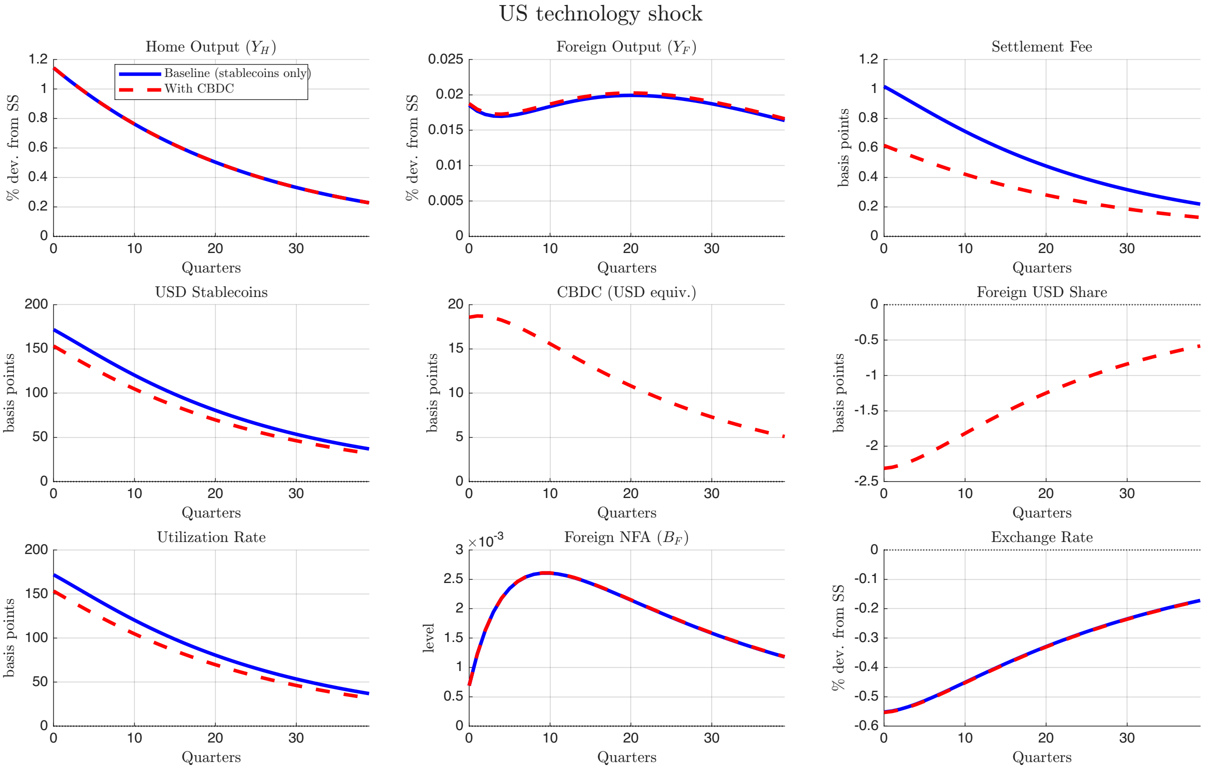

Figure 5: Impulse responses to US productivity shock.

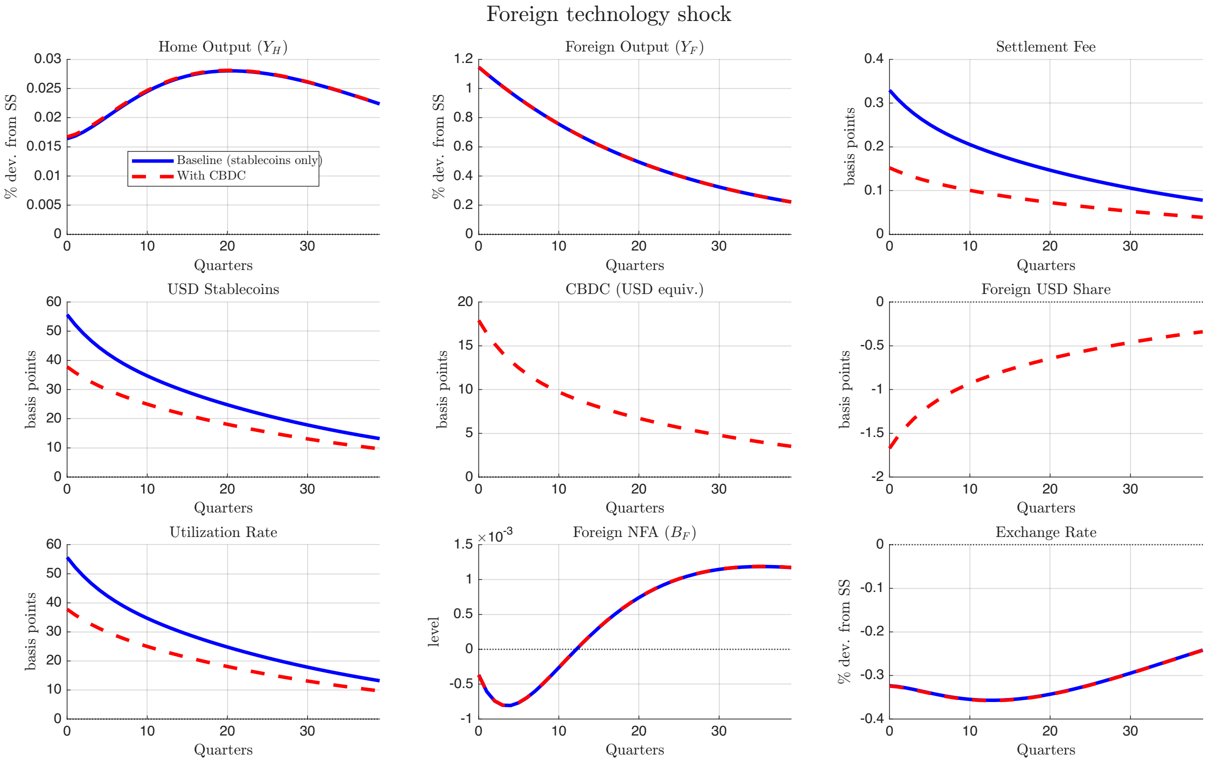

Figure 6: Impulse responses to foreign productivity shock.