Stablecoin Architecture, Digital Dollarization, and the New Triffin Dilemma*

Silvia Sgherri

George Washinton University

This version: March 14, 2026

Abstract

Private stablecoins backed by sovereign bond portfolios have emerged as globally circulating reserve-like assets whose design reallocates currency exposure across segmented international asset markets. This paper develops a three-country New Keynesian model with micro-founded liquidity services and portfolio-balance frictions to study how stablecoin architecture shapes international shock transmission under imperfect asset substitutability. The analysis separates three design dimensions of stablecoins—denomination, unit of account, and backing composition—that are typically treated as equivalent. Each dimension operates through a distinct equilibrium channel. Concentrated denomination transmits peripheral risk-off shocks into advanced-economy exchange rates; diversification eliminates this externality exactly. The unit-of-account basket governs purchasing-power insurance for peripheral households and has no effect on any bilateral exchange rate. Backing misalignment creates procyclical equity destruction; alignment insulates the issuer's balance sheet entirely. Welfare comparisons reveal a sharp quantitative hierarchy across channels. Variation induced by the unit-of-account channel is up to 190 times larger than that generated by denomination choices, reflecting the stablecoin's role as insurance against peripheral currency risk. However, the welfare-maximizing basket differs systematically across countries: the optimal basket for the periphery conflicts with that of the advanced-economy core, and the conflict increases with adoption. This is the *new Triffin Dilemma* in a digitally dollarized world. Crucially, denomination and basket composition operate through orthogonal mechanisms: financial stability regulation and basket governance are thus separable instruments—a property the original Triffin dilemma lacked and that makes the new one institutionally tractable.

JEL: F31, F33, F41, E42, E44, G23

Keywords: Stablecoins, Portfolio Balance, International Monetary System, Reserve Asset Design, Triffin Dilemma, Digital Economy

*I thank seminar and conference participants for valuable comments. All errors are my own.

1 Introduction

When Robert Triffin warned in 1960 that the dollar-based monetary system contained the seeds of its own destruction, the instrument of instability was a sovereign balance sheet. Today it is a smart contract. The logic is the same: the world demands a safe and liquid reserve asset, and providing it redistributes risks in ways that no single issuer can resolve. Private stablecoins have recreated the Triffin structure at scale, outside any multilateral governance, and with a speed the original architects of Bretton Woods could not have imagined.

Dollar-pegged stablecoins—holding nearly $150 billion in US Treasury securities by 2024—have become a shadow reserve system, intermediating global dollar-liquidity demand outside any multilateral governance (Gorton and Zhang, 2023; Catalini et al., 2022). They are used overwhelmingly in emerging markets and developing economies, precisely where domestic currency risk is highest. The substitution of domestic currencies with foreign-denominated stablecoins is commonly referred to as 'digital dollarization', reflecting the empirical dominance of dollar-pegged stablecoins. The demand is genuine: households in high-risk peripheral economies hold stablecoins for the same reason they hold dollars—to shelter purchasing power when their own currency depreciates. In this sense a dollar-pegged stablecoin system is working exactly as designed.

The difficulty is that satisfying the demand for a safe and liquid reserve asset through a concentrated dollar-denominated stablecoin architecture generates a macroeconomic externality for the advanced-economy core. Peripheral risk-off shocks transmit into advanced-economy bilateral exchange rates through the stablecoin portfolio-balance channel, creating contagion that is entirely absent in the no-stablecoin world. The conditions that maximize peripheral welfare simultaneously maximize welfare costs for the core. The conflict widens with every unit of adoption. This is the New Triffin Dilemma in a digitally dollarized world. Unlike its twentieth-century predecessorwhere the reserve asset was supplied through the balance sheet of a sovereign issuerthe new dilemma arises in a system where shadow reserve currency assets are created by private digital intermediaries but backed by sovereign debt.

This paper provides its formal foundations within a structural theory of stablecoin architecture in a multi-country digital monetary system. The central insight is that stablecoin design operates along three distinct marginsdenomination, unit of account, and backing compositioneach entering different equilibrium conditions and generating different international spillovers. Distinguishing these margins reveals both how stablecoins transmit shocks across countries and why their welfare implications differ sharply across core and periphery.

To formalize this environment, we develop a three-country New Keynesian model cali-

brated to a multipolar monetary world.[1] Home (H) and Foreign (F) are advanced-economy currency areas, while the Rest of the World (R) is a high-yield peripheral economy whose assets carry a risk premium. A single global stablecoin circulates across all three economies; its denomination vector, unit-of-account basket, backing composition, and adoption scale are treated as design parameters. Three independent bilateral exchange rates—H-F, H-R, F-R—are each governed by separate UIP conditions with portfolio-balance premia, capturing the full multipolar structure.

We unbundle three design margins that the existing literature has largely treated as equivalent. The *denomination vector* θ^d governs which currencies the stablecoin's liabilities are expressed in; it determines how fluctuations in outstanding supply translate into net currency exposures across segmented asset markets. The *unit-of-account basket* θ^u defines the price index used to compute real purchasing power; it governs the stablecoin's insurance value relative to domestic money across user jurisdictions. The *backing composition* θ^a specifies the asset portfolio held in reserve; it determines how exchange rate movements affect issuer equity. These three vectors operate through entirely different general equilibrium equations, and conflating them leads directly to incorrect predictions about which designs stabilize or destabilize the system.

The paper establishes five results.

Proposition 1 (Denomination Neutrality). Under symmetric diversified denomination—equal weights on H, F, and R—the stablecoin-induced component of every bilateral UIP premium is identically zero. Not approximately zero. Not zero on average. Zero at every shock realization, every period, every level of adoption. This is a structural result, not a calibration artifact.

Proposition 2 (Cross-Market Contagion and Its Elimination). Under concentrated H-denomination, a peripheral risk-off shock that in the no-stablecoin world would leave the H-F bilateral undisturbed transmits directly into ψ_{hf} through the denomination channel. Denomination diversification severs this transmission belt completely (Figure 1).

Proposition 3 (Backing Alignment as Balance-Sheet Insurance). Aligning backing weights with denomination weights hedges the issuer's equity against exchange rate movements at first order. Misalignment creates procyclical equity destruction that can trigger a doom loop when equity-feedback premia are active (Figure 4).

Proposition 4 (Unit-of-Account Orthogonality). The unit-of-account (UoA) basket operates exclusively through relative money demand, shifting domestic IS curves without

[1] The term 'multipolar' here refers to a world in which private stablecoins extend the reach of dominant-currency denomination across the periphery, intensifying rather than resolving the existing reserve hierarchy's structural tensions.

affecting any bilateral exchange rate or issuer equity. It is structurally orthogonal to both the denomination and backing channels.

Proposition 5 (The New Triffin Dilemma). Welfare analysis reveals a strict hierarchy of design margins: the unit-of-account channel generates welfare variation up to 190 times larger than denomination along the risky-currency simplex edge (Table 2). The stablecoin functions primarily as peripheral currency insurance. But the welfare-maximising UoA for R households ($x_R^* \approx 0.12$, heavy in own currency) conflicts irreconcilably with that of the advanced-economy bloc ($x_{HF}^* \approx 0.91$, near R-exclusion), with world welfare at an interior solution ($x_W^* \approx 0.65$) satisfying neither. The conflict widens monotonically with adoption.

The key feature that distinguishes the New Triffin Dilemma from its predecessor is *separability*. Proposition 4 implies that the financial-stability problem (denomination regulation) and the distributional problem (basket governance) rely on orthogonal policy instruments and can therefore be assigned to different institutions. Denomination regulation—requiring the knife-edge design of symmetric denomination and aligned backingbelongs with financial-stability authorities such as the FSB or BIS. Basket governance—determining the world-welfare-optimal unit-of-account composition—requires a multilateral institution representing both core and periphery, for which the IMF is the natural candidate. The original Triffin dilemma lacked this institutional separability, which is precisely why it proved insoluble within the Bretton Woods architecture. The new dilemma, by contrast, has a solution—even if the political economy of achieving it remains difficult.

Relation to the Literature

Three debates in the literature prepare the ground for this paper, and its contribution to each is sharpened by the five results above.

The *multi-country New Keynesian literature* has mapped with precision how monetary policy, exchange rates, and external imbalances propagate across borders (Clarida et al., 2002; Galí and Monacelli, 2005; Obstfeld and Rogoff, 2002; Corsetti et al., 2010; Devereux and Engel, 2003; Engel, 2011). Our contribution is to embed a globally circulating private asset whose denomination vector shifts UIP equilibria, introducing a contagion channel absent from the standard two-country framework. The three-country structure here is not a generalization for its own sake: the transmission from a peripheral risk-off shock into advanced-economy bilaterals through the stablecoin portfolio-balance channel is invisible in any two-country model. It requires a third country to exist, and a stablecoin to connect them.

The *safe-asset and reserve-currency literature* has established why the world economy

generates chronic excess demand for safe stores of value, why that demand concentrates on a single reserve currency, and why the resulting system is fragile (Caballero et al., 2008, 2017; Farhi and Maggiori, 2018; He et al., 2019; Maggiori et al., 2020; Gopinath and Stein, 2021; Gourinchas and Rey, 2007; Gourinchas et al., 2017). This strand of the literature asks who supplies the safe asset and at what cost to the issuer. Our paper asks a different question: what happens when the safe asset is supplied by a private intermediary whose design choices denomination, backing, unit of account are made for profit rather than systemic stability? The answer is that the welfare hierarchy inverts. The reserve-quality debate, which focuses on backing composition, turns out to address a second-order lever: the unit-of-account channel generates welfare variation up to 190 times larger. The instrument that regulators are currently debating is not the one that matters most.

The *stablecoin and digital-currency literature* has grown rapidly around questions of monetary sovereignty, regulatory arbitrage, and the competitive implications of private digital money (Brunnermeier et al., 2019; Benigno et al., 2023; Catalini et al., 2022; Gorton and Zhang, 2023; Auer et al., 2021; Mancini Griffoli et al., 2018; Adrian and Mancini Griffoli, 2019). This work has been largely institutional and partial-equilibrium in character, which is appropriate for first-pass policy analysis but leaves the general equilibrium externalities uncharacterized. Our contribution is a framework with exact analytical results: denomination neutrality holds precisely at the knife-edge, not approximately; the unit of account and denomination channels are structurally orthogonal, not merely empirically separable; and the New Triffin Dilemma has a formal welfare statement that identifies the conflict, quantifies it, and locates it in the adoption trajectory. The dilemma is not a metaphor. It is a proposition.

The paper proceeds as follows. Section 2 develops the model. Section 3 derives the main propositions, closing with a foreshadowing of the dilemma structure. Section 4 presents the quantitative analysis in five steps: the core contagion result, the transmission mechanism, the real consequences, the balance-sheet doom loop, and the welfare gradient. Section 5 develops the two-instrument regulatory framework. Section 6 concludes. Appendices contain the full log-linearized system, proofs, and Blanchard–Kahn conditions.

2 Model

2.1 Country Structure and Exchange Rates

The world consists of three countries: Home (H), Foreign (F), and Rest of World (R). H and F are advanced economies with deep liquid bond markets; R is a peripheral emerging

market whose assets carry a risk premium. All variables are log-deviations from a symmetric non-stochastic steady state. Three bilateral nominal exchange rates are independently determined: e_{hf} (H per F), e_{hr} (H per R), and e_{fr} (F per R). Each has its own independent UIP condition, capturing the full multipolar structure without any cross-rate identity imposed as a constraint.

2.2 Households

The representative household in country $j \in \{H, F, R\}$ maximizes:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_{j,t}^{1-\sigma}}{1-\sigma} - \frac{\varphi_l}{2}(C_{j,t} - L_{j,t})^2 \right], \tag{1}$$

where $\beta \in (0,1)$, $\sigma > 0$, and $\varphi_l > 0$ is the liquidity wedge strength. Liquidity services aggregate domestic money and stablecoins:

$$l_j = (1-\omega)\, mb_j + \omega\, s_j, \tag{2}$$

with $\omega \in [0,1]$ the adoption share. The log-linearized Euler equation is:

$$(\sigma + \varphi_l)\, c_j = (\sigma + \varphi_l)\, \mathbb{E}_t c_j^{+1} - \varphi_l\, \mathbb{E}_t(l_j^{+1} - l_j) - (i_j - \mathbb{E}_t \pi_j^{+1}). \tag{3}$$

The term $\varphi_l\, \mathbb{E}_t(l_j^{+1} - l_j)$ is the IS-curve channel through which stablecoin quantity and unit-of-account dynamics affect output. Standard money demand gives $mb_j = -\eta\, i_j$. The relative demand for stablecoins versus money is governed by the purchasing-power wedge:

$$mb_j - s_j = \eta(\pi_j^{\text{sc}} - \pi_j), \tag{4}$$

where π_j^{sc} is stablecoin basket inflation in country j's numeraire. This is the primary channel through which the unit-of-account parameter affects the real economy: when basket inflation diverges from domestic CPI, households reallocate between money and stablecoins.

2.3 Firms, Supply Side, and Monetary Policy

Monopolistically competitive firms set prices under Calvo contracts (Calvo, 1983). The New Keynesian Phillips Curve is:

$$\pi_{jd} = \beta\, \mathbb{E}_t \pi_{jd}^{+1} + \kappa\, y_j. \tag{5}$$

CPI inflation aggregates domestic and imported components with symmetric import shares $\gamma/2$ from each trading partner:

$$\pi_H = (1-\gamma)\,\pi_{Hd} + \tfrac{\gamma}{2}(\pi_{Fd} + \dot{e}_{hf}) + \tfrac{\gamma}{2}(\pi_{Rd} + \dot{e}_{hr}), \tag{6}$$

$$\pi_F = (1-\gamma)\,\pi_{Fd} + \tfrac{\gamma}{2}(\pi_{Hd} - \dot{e}_{hf}) + \tfrac{\gamma}{2}(\pi_{Rd} + \dot{e}_{fr}), \tag{7}$$

$$\pi_R = (1-\gamma)\,\pi_{Rd} + \tfrac{\gamma}{2}(\pi_{Hd} - \dot{e}_{hr}) + \tfrac{\gamma}{2}(\pi_{Fd} - \dot{e}_{fr}). \tag{8}$$

Each central bank follows a Taylor rule with smoothing ρ_i, inflation response $\phi_\pi > 1$, and output response $\phi_y > 0$.

2.4 Exchange Rate Determination and the Risk Premium

Exchange rates are determined by independent UIP conditions with portfolio-balance risk premia (Gabaix and Maggiori, 2015; Obstfeld et al., 2009):

$$e_{hf} = \mathbb{E}_t e_{hf}^{+1} + i_H - i_F - \psi_{hf}, \tag{9}$$

$$e_{hr} = \mathbb{E}_t e_{hr}^{+1} + i_H - i_R - \psi_{hr}, \tag{10}$$

$$e_{fr} = \mathbb{E}_t e_{fr}^{+1} + i_F - i_R - \psi_{fr}. \tag{11}$$

Each risk premium combines an NFA term, a stablecoin-induced term, an equity-feedback term, and an exogenous shock:

$$\psi_{hf} = \psi_b\, b_{hf} + \psi_{hf}^{\text{sc}} - \chi_E\, E^{\text{sc}}, \tag{12}$$

$$\psi_{hr} = \psi_b\, b_{hr} + \psi_{hr}^{\text{sc}} - \chi_E\, E^{\text{sc}} + \nu, \tag{13}$$

$$\psi_{fr} = \psi_b\, b_{fr} + \psi_{fr}^{\text{sc}} - \chi_E\, E^{\text{sc}} + \nu. \tag{14}$$

The global risk-off shock ν enters ψ_{hr} and ψ_{fr} but *not* ψ_{hf}. This captures a multipolar world in which H and F are symmetric safe-haven areas: a risk-off episode elevates the cost of R borrowing against both advanced economies, but leaves the H-F bilateral undisturbed in the absence of stablecoin-induced spillovers.

The equity-feedback coefficient $\chi_E \geq 0$ governs the extent to which issuer equity dampens or amplifies the overall premium. When the equity buffer is positive, $-\chi_E E^{\text{sc}} < 0$ is stabilising. When the buffer is depleted and equity turns negative, the same term reverses sign and amplifies—the exhaustion doom loop formalized in Corollary 1.

2.5 The Stablecoin: Design Parameters and Equilibrium Channels

Design vectors. The stablecoin is characterized by three probability simplices over $\{H, F, R\}$: the denomination vector $\theta^d = (\theta^d_H, \theta^d_F, \theta^d_R)$; the backing vector $\theta^a = (\theta^a_H, \theta^a_F, \theta^a_R)$; and the unit-of-account vector $\theta^u = (\theta^u_H, \theta^u_F, \theta^u_R)$. These three vectors are in general independent and operate through entirely distinct equilibrium equations.

Denomination and the UIP wedge. The stablecoin-induced UIP wedges are:

$$\psi^{\text{sc}}_{hf} = \psi_b(\theta^d_H - \theta^d_F)\bar{s}(s - \zeta_E E^{\text{sc}}), \tag{15}$$

$$\psi^{\text{sc}}_{hr} = \psi_b(\theta^d_H - \theta^d_R)\bar{s}(s - \zeta_E E^{\text{sc}}), \tag{16}$$

$$\psi^{\text{sc}}_{fr} = \psi_b(\theta^d_F - \theta^d_R)\bar{s}(s - \zeta_E E^{\text{sc}}). \tag{17}$$

The concentration differences $(\theta^d_j - \theta^d_k)$ are the structural coefficients that translate stablecoin quantity fluctuations into bilateral UIP wedges. Setting $\theta^d_H = \theta^d_F = \theta^d_R = 1/3$ sets all three coefficients to zero simultaneously.

Unit of account and money demand. Stablecoin basket inflation in H numeraire is:

$$\pi^{\text{sc}} = \theta^u_H \pi_{Hd} + \theta^u_F(\pi_{Fd} + \dot{e}_{hf}) + \theta^u_R(\pi_{Rd} + \dot{e}_{hr}). \tag{18}$$

This feeds into the MB-SC demand wedge (4) and through it the IS equations. It has no direct pathway to UIP premia.

Stablecoin aggregation. $s = \lambda_s s_H + \frac{1-\lambda_s}{2} s_F + \frac{1-\lambda_s}{2} s_R$.

2.6 The Issuer's Balance Sheet

The issuer holds a backing portfolio with nominal return:

$$r^A = \theta^a_H i_H + \theta^a_F(i_F + \mathbb{E}_t \dot{e}^{+1}_{hf}) + \theta^a_R(i_R + \mathbb{E}_t \dot{e}^{+1}_{hr}). \tag{19}$$

Real profit is $\pi_{\text{prof}} = \bar{s}(r^A - \pi^{\text{sc}})$. Equity evolves as:

$$E^{\text{sc}}_t = \beta^{-1} E^{\text{sc}}_{t-1} + \pi_{\text{prof}} + \text{val}_q + \text{val}_{fx} - \phi_{\text{div}} E^{\text{sc}}_{t-1}, \tag{20}$$

where $\text{val}_q = \bar{s} \sum_j \theta^a_j \Delta q_j$ and the FX valuation term is:

$$\text{val}_{fx} = \bar{s}\big[(\theta^a_F - \theta^d_F)\dot{e}_{hf} + (\theta^a_R - \theta^d_R)\dot{e}_{hr}\big]. \tag{21}$$

Under backing-denomination alignment $\theta_j^a = \theta_j^d$ for all j, equation (21) vanishes identically: equity is exchange-rate insulated. Under misalignment with $\theta_R^a < \theta_R^d$ — the issuer underweights R assets relative to R-denomination liabilities—R depreciation ($\dot{e}_{hr} < 0$) generates $\text{val}_{fx} < 0$: equity is destroyed procyclically.

Dividends are rebated to households proportionally to SC holdings: $T_H = \lambda_s \phi_{\text{div}} E_{t-1}^{\text{sc}}$ and similarly for F and R.

2.7 External Balance and Goods Markets

Bilateral NFAs evolve as $b_{jk,t} = \beta^{-1} b_{jk,t-1} + nx_{jk,t}$ with net exports $nx_{jk} = \frac{\gamma}{2}(c_k - c_j) + \frac{\gamma}{2}\text{tot}_{jk}$. Goods market clearing conditions yield $y_j = c_j + nx_j - T_j$ for each country.

3 Analytical Results

3.1 Denomination Neutrality

Proposition 1 (Denomination Neutrality)**.** *Under symmetric diversified denomination, $\theta_H^d = \theta_F^d = \theta_R^d = 1/3$, the stablecoin-induced UIP wedges are identically zero: $\psi_{hf}^{\text{sc}} = \psi_{hr}^{\text{sc}} = \psi_{fr}^{\text{sc}} = 0$ for all t, all shock realizations, and all levels of adoption ω.*

Proof. See Appendix B. □

The result is structural, not a calibration artefact. Denomination diversification eliminates the contagion channel exactly—not approximately, not on average, but identically at every shock realization and every adoption level. Put differently: a stablecoin can grow arbitrarily large under the knife-edge design without ever introducing a single basis point of denomination-channel UIP distortion.

3.2 Cross-Market Contagion and Its Elimination

Proposition 2 (Cross-Market Insulation)**.** *Under symmetric diversified denomination, changes in s induced by shocks to ψ_{hr} and ψ_{fr} have no first-order effect on ψ_{hf} or e_{hf}. Under concentrated denomination ($\theta_H^d > \theta_F^d$), such changes enter ψ_{hf}^{sc} with a positive coefficient, transmitting the peripheral shock directly into the H-F market.*

Proof. See Appendix B. □

The mechanism is a portfolio-balance effect on the liability side. Dollar-denominated stablecoins create a net long position in H bonds in segmented markets; any shock that

moves aggregate demand s shifts ψ_{hf} proportionally to denomination concentration, even if the shock originates entirely in H-R or F-R markets. The three-country structure is essential here: this channel is invisible in any two-country model.

3.3 Backing Alignment as Balance-Sheet Insurance

Proposition 3 (Backing-Denomination Alignment)**.** *Under alignment $\theta_j^a = \theta_j^d$ for all j, $\text{val}_{fx} \equiv 0$ for all exchange rate realisations. Under misalignment with $\theta_R^a > \theta_R^d$ and R depreciation ($\dot{e}_{hr} < 0$), $\text{val}_{fx} < 0$: equity is procyclically destroyed when the shock is largest.*

Proof. See Appendix B. □

Corollary 1 (The Exhaustion Doom Loop)**.** *Under alignment and $\chi_E > 0$, equity stability ensures $-\chi_E E^{\text{sc}} \approx 0$ throughout, providing no amplification. Under misalignment with $\chi_E > 0$, equity destruction on impact causes $-\chi_E E^{\text{sc}} > 0$, raising the premium and amplifying the initial shock. Once the equity buffer is exhausted, the feedback reverses sign: what was stabilising on impact becomes destabilising from period 3 onward. Architecture and capital requirements are complements for the wrong reason under misalignment—requiring capital from a misaligned issuer partially offsets an amplification that aligned architecture would have prevented entirely.*

3.4 Unit-of-Account Orthogonality

Proposition 4 (UoA-Denomination Orthogonality)**.** *For any fixed θ^d and θ^a, the equilibrium paths of ψ_{hf}^{sc}, ψ_{hr}^{sc}, ψ_{fr}^{sc}, val_{fx}, and all bilateral exchange rates are independent of θ^u at first order. The UoA channel operates exclusively through the IS equations via the demand wedge $mb_j - s_j$ and liquidity services l_j.*

Proof. See Appendix B. □

3.5 The Knife-Edge Design and a Preview of the Dilemma

Propositions 1 and 3 jointly define the *knife-edge*: symmetric diversified denomination combined with denomination-aligned backing. Under this design all three stablecoin-induced UIP wedges are identically zero and issuer equity is exchange-rate insulated. The unit-of-account parameter remains free.

Corollary 2 (Uniqueness)**.** *The knife-edge is the unique architecture satisfying $\psi_{hf}^{\text{sc}} = \psi_{hr}^{\text{sc}} = \psi_{fr}^{\text{sc}} = 0$ and $\text{val}_{fx} = 0$ simultaneously for all exchange rate realizations.*

These propositions solve the financial stability problem. But they do not solve the distributional problem—and it is the distributional problem that constitutes the New Triffin Dilemma. Proposition 4 implies that the unit-of-account basket can be chosen freely without disturbing any of the neutrality results above. It also implies that the welfare consequences of the UoA choice are entirely independent of whether the issuer has adopted the knife-edge. The welfare analysis of Section 4.6 will show that the UoA dimension generates welfare variation an order of magnitude larger than denomination—precisely because the stablecoin functions primarily as insurance against peripheral currency risk, not as a financial market instrument. And it will show that the welfare-maximising basket for R households is incompatible with the welfare-maximising basket for the advanced-economy core. No design choice resolves this conflict from within. It requires governance from without.

4 Quantitative Analysis

We proceed in five steps: the core contagion result (Section 4.2); the transmission mechanism (Section 4.3); the real consequences (Section 4.4); the balance-sheet doom loop (Section 4.5); and the welfare gradient that culminates in the New Triffin Dilemma (Section 4.6).

4.1 Calibration and Regime Definitions

Calibration. Parameters are set at quarterly frequency following conventional values in the open-economy DSGE literature. Table 1 summarizes. The portfolio-balance coefficient $\psi_b = 0.60$ is higher than the conventional 0.30 used in two-country models for three reasons: a three-currency world has more bilateral pairs over which exposure is priced; the denomination-channel amplification scales with $\psi_b \cdot \bar{s}$; and higher ψ_b matches emerging-market evidence on portfolio-rebalancing effects (Gabaix and Maggiori, 2015). The payout rate $\phi_{\text{div}} = 1$ sets the baseline as full pass-through: all issuer profits and valuation gains return to households each period, so equity is a pure accounting identity with no causal role in equilibrium dynamics. This provides the cleanest separation between the denomination and backing channels and any buffer effects; Section 4.5 activates buffers by switching to $\phi_{\text{div}} = 0.5$, $\chi_E = 0.025$, $\zeta_E = 0.05$. The discount factor $\beta = 0.99$ implies an annual real rate of approximately 4 percent, consistent with advanced-economy calibrations; R's higher effective cost of capital is captured through the risk premium ν rather than a country-specific discount factor. The adoption share $\omega = 0.50$ is a scenario calibration; the empirically current share is approximately 1.5 percent of broad money, implying negligible welfare differences across regimes. The scenarios illustrate the consequences of the growth trajectory.

Table 1: Baseline Calibration

Parameter	Description	Value	Source / Target
β	Discount factor	0.99	$r \approx 4\%$ ann.
σ	Risk aversion	3.0	DSGE standard
η	Interest semi-elasticity	2.0	Money demand
κ	NKPC slope	0.12	4-quarter Calvo
γ	Import share	0.30	Trade data
φ_l	Liquidity wedge	0.50	SC adoption impact
ρ_i	Taylor smoothing	0.80	Standard
ϕ_π	Taylor inflation	1.50	Taylor principle
ϕ_y	Taylor output	1.00	Standard
ψ_b	Portfolio-balance coeff.	0.60	Three-currency calibration
ρ_ν	Risk premium persistence	0.80	VIX persistence
ω	SC adoption share	0.50	Scenario
ϕ_{div}	Issuer payout rate	1.00	Full pass-through baseline
λ_s	H share of SC holdings	1/3	Symmetric
χ_E	Equity-to-premium coeff.	1.00	Baseline
ζ_E	Equity-to-exposure coeff.	0.00	Baseline

Regime definitions. Four design regimes span the relevant corners of the architecture space. *Regime 1* ($\omega = 0$) is the no-stablecoin benchmark. *Regime 2* sets $\theta^d = \theta^a = \theta^u = (1, 0, 0)$: the dollar stablecoin archetype—concentrated H denomination, H-backed, H unit of account. Denomination contagion is maximal; balance sheet is exchange-rate insulated. *Regime 3* sets $\theta^d = \theta^u = (1, 0, 0)$ and $\theta^a = (0, 0, 1)$: same denomination and UoA as Regime 2, but fully R-backed. The issuer runs a carry trade: H-currency liabilities funded by R assets. When R depreciates, $\text{val}_{fx} = \bar{s}(\theta^a_R - \theta^d_R)\dot{e}_{hr} = \bar{s} \cdot (1 - 0) \cdot \dot{e}_{hr} < 0$: equity is destroyed procyclically. *Regime 4* sets $\theta^d = \theta^a = \theta^u = (1/3, 1/3, 1/3)$: the knife-edge. All denomination-channel and backing-channel externalities are identically zero.

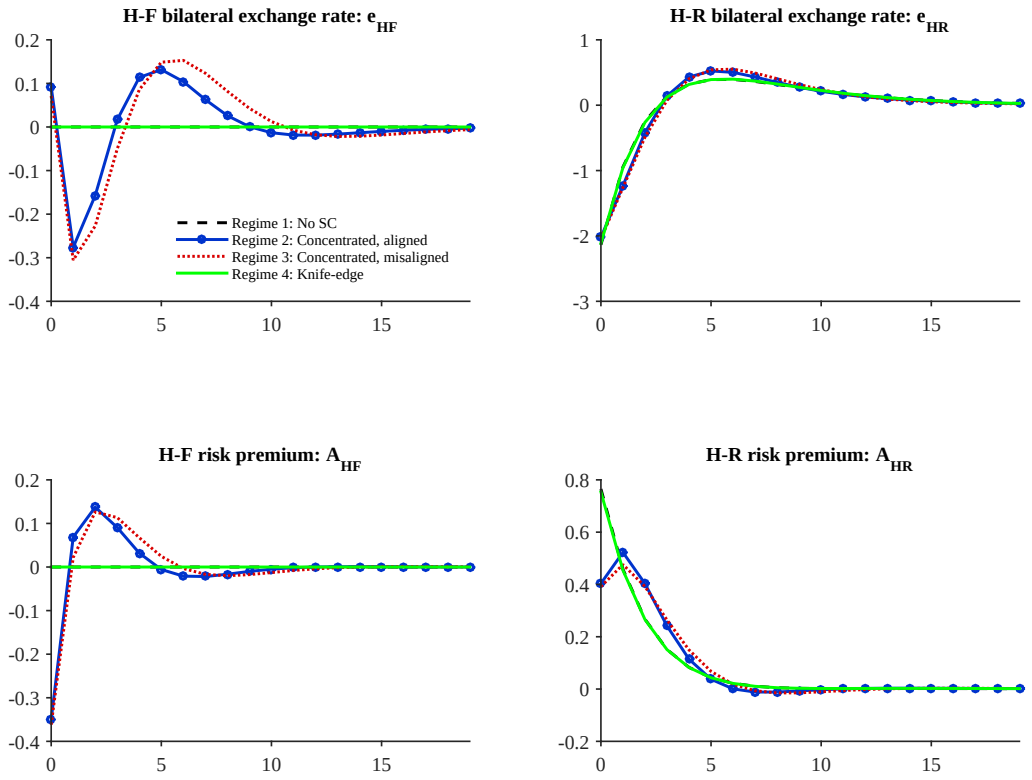
Regimes 2 and 3 share identical denomination, so any difference between them isolates the backing-misalignment channel. Regime 4 versus Regime 1 confirms denomination neutrality: the knife-edge should track the no-stablecoin benchmark in all UIP-affected variables.

4.2 The Core Result: Denomination Neutrality and Contagion

Figure 1 is the paper's central result. Read the right column first: all four regimes produce nearly identical e_{hr} and ψ_{hr} responses—R depreciates sharply against H, the premium decays at rate $\rho_\nu = 0.80$. This is the directly shocked bilateral, and architecture does not alter it materially.

Now read the left column. This is the test of Proposition 2. In the no-stablecoin world

Figure 1: Denomination Neutrality and Contagion



Note: Impulse responses to a one-standard-deviation global risk-off shock ν, entering ψ_{hr} and ψ_{fr} simultaneously. Top: H-F bilateral exchange rate e_{hf} (left) and H-R rate e_{hr} (right). Bottom: H-F risk premium ψ_{hf} (left) and H-R premium ψ_{hr} (right). Regime 1 (black dashed): no-stablecoin benchmark. Regime 2 (blue with marker): concentrated H-denomination, aligned backing. Regime 3 (red dotted): concentrated H-denomination, R-backed. Regime 4 (green solid): knife-edge.

(black dashed), a purely R-specific shock leaves the H-F bilateral completely undisturbed—flat at zero throughout. Under Regimes 2 and 3, a ψ_{hf} spike of approximately 0.15 appears on impact. A shock with no fundamental connection to the H-F bilateral is generating advanced-economy exchange rate volatility through the stablecoin denomination channel. Under Regime 4, the solid green line lies exactly on the black dashed line throughout. The knife-edge design generates zero H-F response to a purely R shock—at every horizon, not just on average.

The Regime 2 and Regime 3 H-F responses are nearly identical, as theory requires: both have $\theta^d = (1, 0, 0)$, so ψ^{sc}_{hf} is the same function of s. Denomination determines contagion; backing does not. This is Proposition 2 in one panel.

4.3 The Transmission Mechanism

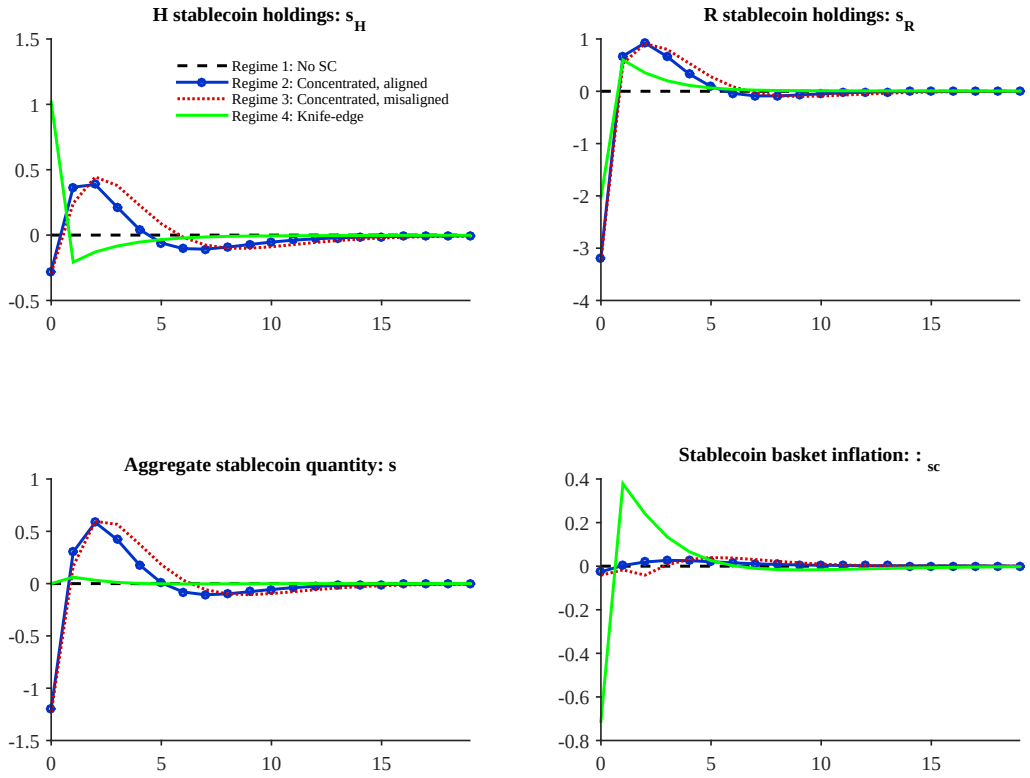
The four panels of Figure 2 trace the shock from its origin in R currency risk through to its resolution in portfolio stocks.

The s_R collapse (top-right) is the most violent movement in the figure. On impact, R households briefly surge into stablecoins as a safe haven, producing a spike at period 0. But R's depreciation immediately raises the H-numeraire cost of holding the stablecoin, and by period 1 R households have shed roughly three standard deviations of holdings under Regime 3, and approximately half that under Regime 2. The knife-edge (Regime 4) produces a shallower collapse because the symmetric denomination means the stablecoin retains more purchasing power in R units during the depreciation—the mixed UoA acts as a partial hedge for R households even while the underlying shock is unchanged.

The s_H absorption (top-left) follows with a one-period lag. As R households sell, H and F households buy at the lower price. s_H rises to a peak around period 2 and decays slowly over the subsequent ten periods. Aggregate s (bottom-left) inherits this hump: a large initial contraction as R demand collapses, then a prolonged above-steady-state period as H absorption dominates. It is this aggregate hump that enters the premium equation. Under concentrated denomination, the coefficient is 1 and the hump translates directly into ψ_{hf}. Under the knife-edge it is also present—stablecoin activity is not suppressed—but the coefficient is zero and no premium results.

The basket inflation panel π^{sc} (bottom-right) deserves separate attention because it carries the proof of Proposition 4. Under Regime 4, the symmetric UoA means the basket contains an R-currency component whose H-numeraire value falls sharply when R depreciates: π^{sc} swings from -0.8 on impact to above $+0.4$ in period 2 before mean-reverting. This is the largest single impulse response in the entire figure. Under Regimes 2 and 3 with

Figure 2: The Transmission Chain



Note: Impulse responses to the same risk-off shock as Figure 1. Top-left: H-country stablecoin holdings s_H. Top-right: R-country holdings s_R. Bottom-left: aggregate stablecoin quantity s. Bottom-right: basket inflation π^{sc}. On impact, R households generate a brief safe-haven spike in s_R; by period 1, R depreciation makes the H-denominated stablecoin prohibitively expensive and s_R collapses to its most negative value, with Regime 3 (misaligned backing) showing the deepest reversal. H households absorb the displaced supply with a one-period lag, so s_H and aggregate s follow hump-shaped paths. Under concentrated denomination (Regimes 2 and 3), this aggregate expansion enters $\psi_{hf}^{\text{sc}} = \psi_b(\theta_H^d - \theta_F^d)\bar{s} \cdot s$ with a unit coefficient, delivering the premium spike of Figure 1. The basket inflation panel is the visual proof of UoA orthogonality: under Regime 4 (knife-edge, symmetric basket), π^{sc} collapses to -0.8 on impact as the R-currency weight reprices with the depreciation, then spikes above $+0.4$ as it reverts—yet $\psi_{hf} = 0$ throughout (Figure 1). Under Regimes 2 and 3 with H-only UoA, π^{sc} barely moves. Maximum basket volatility coincides with zero premium contagion. The two dimensions are orthogonal.

pure H-denomination in the UoA, $\pi^{sc} \approx \pi_{Hd}$ and barely moves. The contrast is stark: the regime with violent basket inflation (R4) is exactly the regime with zero ψ_{hf} response, and the regimes with stable basket inflation (R2, R3) are the ones that transmit a 0.15-standard-deviation premium spike into H-F financial conditions. Basket volatility is not the source of contagion; denomination concentration is.

4.4 Liquidity Wedges, Consumption, and External Adjustment

Figure 3 decomposes the real consequences into three margins. The liquidity wedge row (top) isolates the denomination channel most cleanly: gap_H goes negative on impact under Regimes 2 and 3 the premium spike tightens H money demand immediately while under Regime 4 it briefly turns positive, as the symmetric basket preserves purchasing power. gap_R spikes sharply across all regimes regardless of architecture, confirming that R's IS-curve exposure is driven by the shock itself, not by design.

The consumption row (middle) adds the backing dimension. c_H follows the same regime ranking as the liquidity wedge: denomination concentration prolongs and deepens the H trough. The more striking result is in c_R: consumption collapses on impact across all regimes, but Regime 3 diverges upward from period 3 onward, tracing a pronounced recovery hump through period 10. This is the carry-trade dividend R-backed equity gains are paid out as R household income, partially offsetting the depreciation shock. Aligned backing (Regime 2) produces no such hump. Same denomination, same shock, different recovery: the issuer's balance sheet choice determines whether R households receive a transfer or nothing.

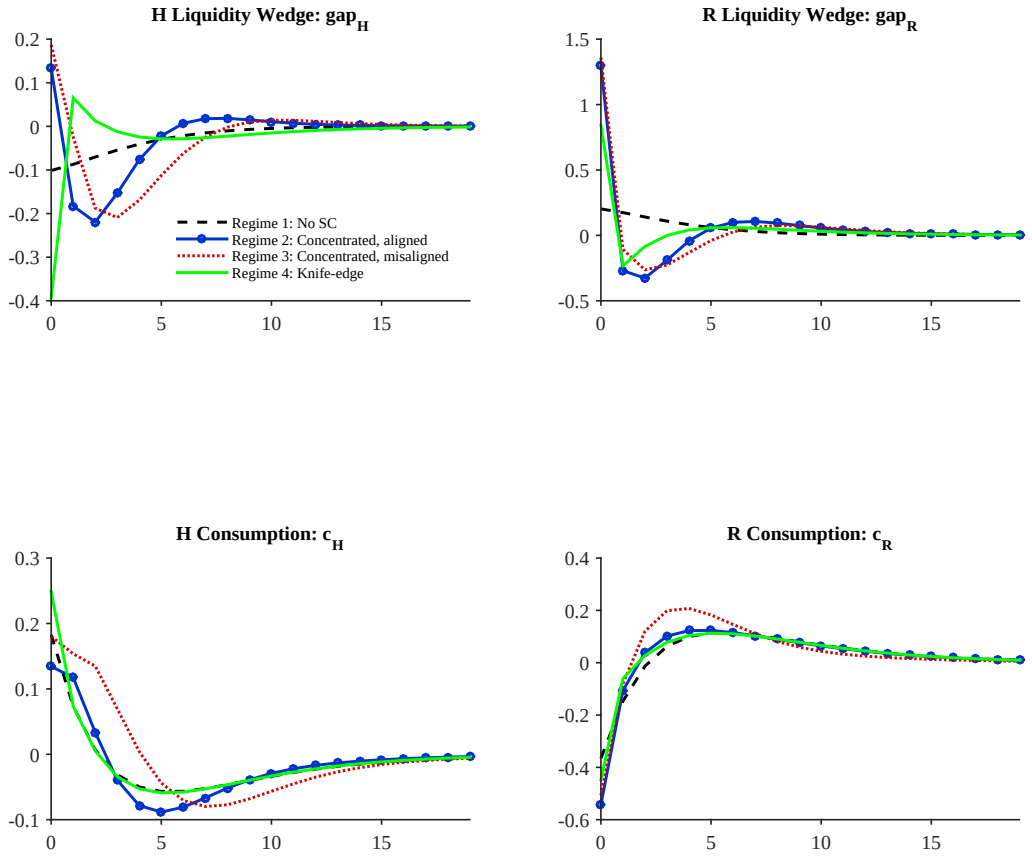
The net export row (bottom) is the control: nx_H and nx_R trajectories are nearly regime-invariant. Stablecoin architecture determines who absorbs the risk internally; it does not alter the external adjustment path in response to a peripheral risk-off shock.

The three-row decomposition thus identifies a clear hierarchy: the liquidity and consumption channels are architecture-sensitive (the denomination determines the sign of gap_H on impact; the backing determines whether c_R recovers through dividend income); the external adjustment channel is architecture-neutral. This hierarchy is precisely what the welfare analysis in Section 4.6 will formalize: the architectural dimensions that matter most for real activity are the same ones that dominate welfare.

4.5 The Balance-Sheet Doom Loop

Three complementary figures characterize the balance-sheet mechanism: the phase diagram (geometry), the second-moment decomposition (magnitude), and the welfare bar chart (cost). Together they establish that architecture dominates buffers, and that the doom loop is a

Figure 3: The Real Consequences

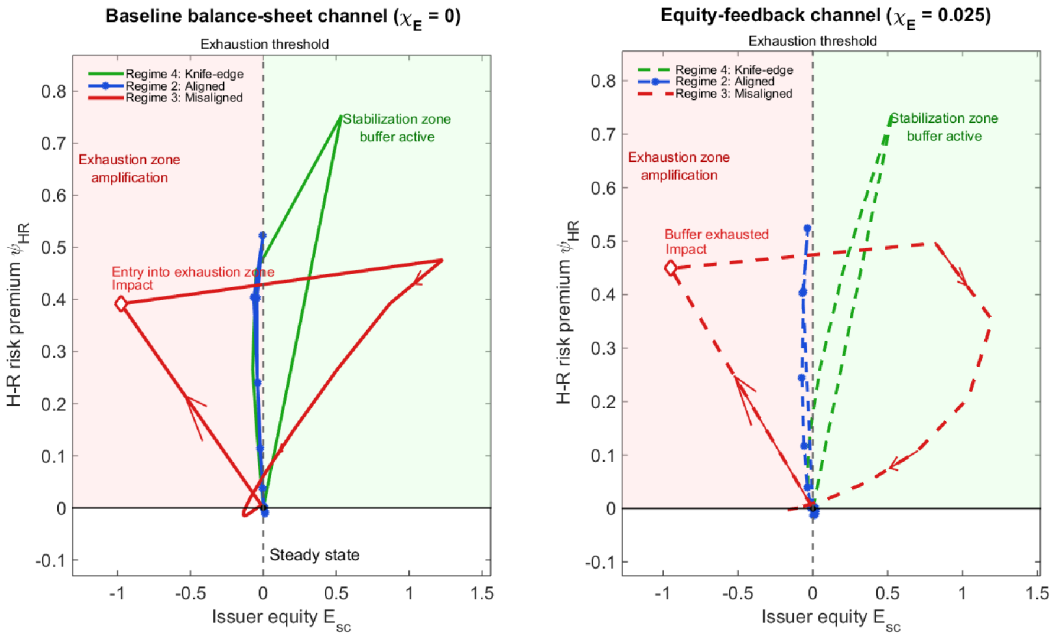


Note: Impulse responses to the same risk-off shock as Figure 1. Top row: H liquidity wedge $\text{gap}_H = c_H - l_H$ (left) and R liquidity wedge gap_R (right). Middle row: H consumption c_H (left) and R consumption c_R (right). Bottom row: H net exports nx_H (left) and R net exports nx_R (right). The top row shows the denomination channel hitting gap_H directly (R4 is positive on impact, R2/R3 are immediately negative) while gap_R spikes positively across all regimes before collapsing. The middle row reveals a consumption asymmetry: c_H is depressed by denomination contagion under R2/R3 while c_R shows a delayed positive hump under Regime 3 from carry-trade dividend payouts. The bottom row shows that external adjustment is largely architecture-invariant: nx_H and nx_R trajectories are nearly identical across all four regimes.

misalignment phenomenon that correctly aligned architecture prevents entirely.

The phase diagram in Figure 4 renders Corollary 1 with full geometric clarity. Each panel plots the trajectory of issuer equity E^{sc} (x-axis) against the H-R risk premium ψ_{hr} (y-axis). The dashed vertical line at $E^{\text{sc}} = 0$ is the exhaustion threshold: to its right (green shading) equity is positive and feedback is stabilising; to its left (red shading) equity is negative and the same feedback term amplifies.

Figure 4: The Doom Loop: The Geometry of the Balance Sheet Mechanism



Note: Phase diagram space: issuer equity E^{sc} (x-axis) versus H-R risk premium ψ_{hr} (y-axis). Green shading ($E^{\text{sc}} > 0$): stabilization zone. Red shading ($E^{\text{sc}} < 0$): exhaustion zone, where the equity-feedback term reverses sign. *Left panel* ($\chi_E = 0$, no feedback): Regime 4 (green) traces the largest loop—entirely in the stabilization zone—because knife-edge backing insulates equity from FX movements and safe-haven inflows accumulate. Regime 2 (blue) makes a smaller stabilization-zone loop. Regime 3 (red) enters the exhaustion zone on impact ("Entry into exhaustion zone—Impact") driven purely by val_{fx} destruction; the entire orbit is in the red zone with no feedback required to create the amplification. *Right panel* ($\chi_E = 0.025$, feedback active): Regime 4's stabilization loop expands as positive equity dampens premia. Regime 3's exhaustion loop expands dramatically ("Buffer exhausted—Impact"): negative equity causes feedback to amplify rather than dampen, widening both the equity trough and the premium peak. Buffers improve well-designed architectures and worsen misaligned ones (Corollary 1).

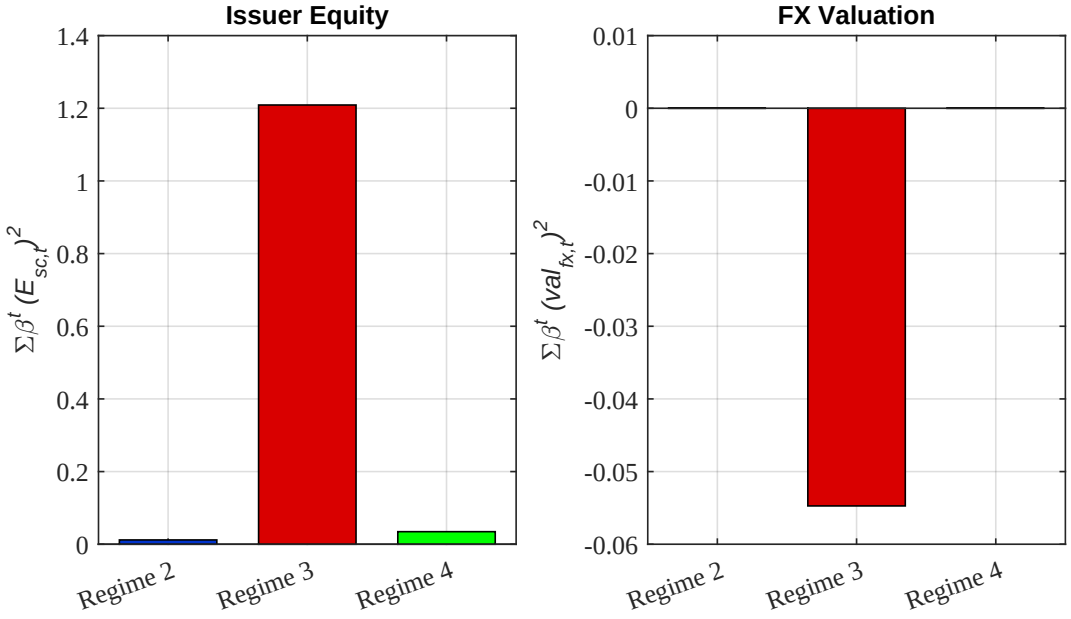
In the left panel ($\chi_E = 0$, no active feedback), the three regime trajectories reveal the architecture mechanism in its purest form. Regime 4 (green, knife-edge) makes the largest loop: the shock generates safe-haven equity inflows that accumulate on the issuer's balance sheet, and because knife-edge backing perfectly aligns assets with liabilities, FX movements

do not erode this equity. The entire Regime 4 orbit traces a wide triangle within the stabilization zone, reaching equity values above +1 before decaying back to the origin. Regime 2 (blue, aligned H-backing) makes a much tighter loop near the origin, also in the stabilization zone. The loop is smaller because aligned H-denomination H-backing leaves little room for valuation gains—the balance sheet is exchange-rate insulated in both directions. Regime 3 (red, R-backed misalignment) is the critical case. On the very first period, the issuer's R-denominated assets are worth less in H numeraire because R has depreciated, destroying equity before any feedback can act. The trajectory crosses the exhaustion threshold at $t = 0$—labelled "Entry into exhaustion zone—Impact" in the figure—and the entire subsequent orbit lies in the red zone at elevated premium levels, tracing a large triangular loop far from the origin. No feedback coefficient is required to create this doom loop. The val_{fx} channel alone, through procyclical equity destruction on impact, generates the exhaustion pathology.

The right panel ($\chi_E = 0.025$, equity feedback active) shows the consequences of adding capital buffers to each architecture. For Regime 4 (green dashed), activating feedback expands the stabilization loop: positive equity now dampens premia, so the trajectory arcs further into the stabilization zone before returning to steady state. Buffers make a well-designed architecture perform even better. For Regime 2 (blue dashed), the effect is small—the loop barely changes because aligned backing keeps equity positive and modest. For Regime 3 (red dashed), the effect is reversal and amplification. The orbit now labelled "Buffer exhausted—Impact" begins already in the exhaustion zone, and because equity is negative, the feedback term $-\chi_E E^{\text{sc}} > 0$ adds to the premium rather than subtracting from it. The Regime 3 loop under active feedback spans a premium range approaching 0.5, compared to roughly 0.4 in the left panel, and the equity trough extends further into negative territory. Buffers designed to stabilize instead amplify when the architecture has already destroyed the equity they depend on. This is the exhaustion doom loop: the instrument that could most help—capital retention—is available only when the architecture has made it most dangerous.

Figure 5 isolates the mechanism through which buffers operate, and the picture contains a striking paradox. The left panel shows the change in equity volatility when buffers are activated. Under Regime 3 (misaligned), the bar reaches approximately +1.2: buffers substantially *increase* equity volatility. This is not a malfunction—it is the geometry of the problem. Retaining equity under Regime 3 means keeping R-denominated assets on the balance sheet longer; those assets continue to generate val_{fx} losses on every subsequent depreciation, amplifying equity swings. Under Regimes 2 and 4, the left bar is near zero: buffers do not move equity volatility when backing is aligned, because there is no exchange-rate exposure to accumulate.

Figure 5: Buffers Effectiveness: The Magnitude of the Balance Sheet Mechanism



Note: Left panel: change in issuer equity discounted second moment $\sum_t \beta^t (E_{sc,t})^2$ from activating buffers, by regime. Right panel: change in FX valuation discounted second moment $\sum_t \beta^t (\text{val}_{fx,t})^2$ from activating buffers. Under Regime 3, buffers sharply increase equity volatility (left bar $\approx +1.2$) while reducing val_{fx} volatility (right bar ≈ -0.055): retained R-denominated equity absorbs valuation shocks but amplifies equity swings. Under Regimes 2 and 4, both bars are near zero—the val_{fx} channel does not exist under aligned backing, so buffers have no mechanism to act through.

The right panel resolves the paradox. It shows the change in val_{fx} volatility from activating buffers. Under Regime 3, the bar is large and negative (≈ -0.055): buffers succeed in reducing the volatility of FX valuation hits by slowing the pace at which R-denominated losses flow through to premium spikes. This is the one channel where buffers genuinely help. Under Regimes 2 and 4, the right- panel bars are exactly zero, because aligned backing means exchange rate movements never hit the balance sheet in the first place.

Together, the two panels state a precise condition for buffer effectiveness: capital requirements can only stabilize a channel that exists. The val_{fx} channel exists only under misaligned backing (Regime 3). Under the architecture most common in practice—H-denomination with H-aligned backing (Regime 2)—neither bar in Figure 5 moves when buffers are activated. The entire Basel-style capital regulation toolkit is inert. Denomination architecture, not capital adequacy, is the operative welfare instrument in the most empirically relevant regime.

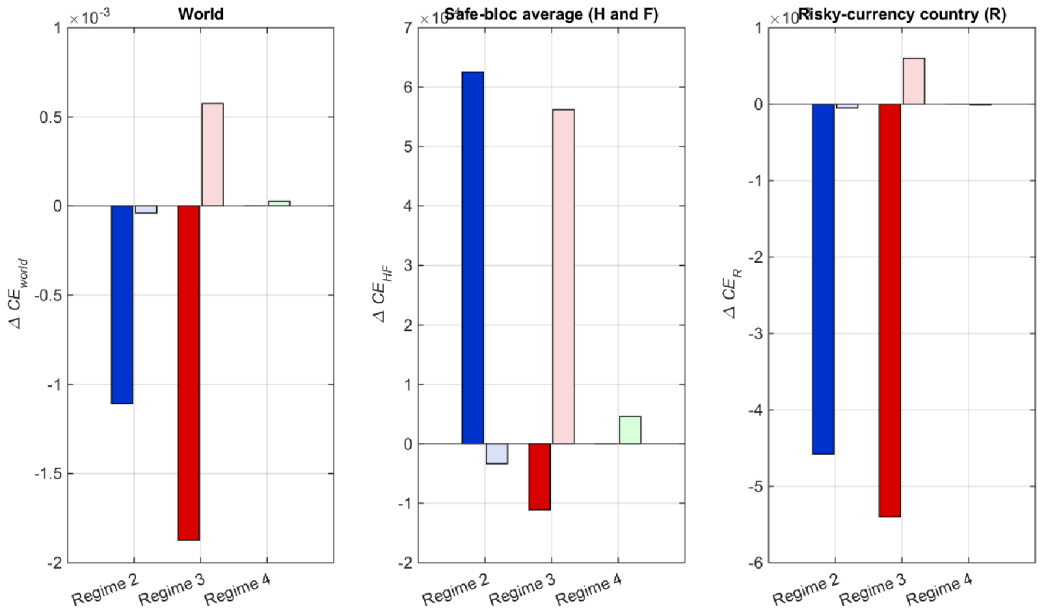
Figure 6 compresses the distributional conflict of the New Triffin Dilemma into three bar charts.

The world panel (left) shows moderate aggregate costs: Regime 2 generates a world welfare loss of approximately -1.1×10^{-3} relative to the knife-edge, and Regime 3 a loss of -1.8×10^{-3}. The pale buffer bars are small at the world level, because the gains buffers deliver to R through the val_{fx} channel are roughly offset by the equity volatility amplification they generate at the issuer level.

The safe-bloc panel (centre) contains the result that would most surprise a standard financial stability analysis. Under Regime 2, the solid blue bar is *positive*—approximately $+6 \times 10^{-4}$. The advanced economies of H and F are *better off* under concentrated H-denomination than under the knife-edge. Why? The denomination channel that creates H-F contagion also creates a safe-haven demand surge for H-currency stablecoins during the shock. H households who hold these claims receive a wealth transfer through the stablecoin premium: the same mechanism that imposes a cost on R is a source of gain for H. Under Regime 3, this safe-bloc gain disappears: misaligned backing destroys issuer equity that H dividend recipients partially own, and the safe-bloc bar turns negative. Buffer effects for the safe bloc are negligible in both regimes—the pale centre bars are tiny—confirming that capital regulation of the issuer is not primarily a tool for protecting the advanced-economy core.

The R-country panel (right) is where the argument lands with full force. The solid blue bar reaches -4.6×10^{-3} under Regime 2 and the solid red bar -5.4×10^{-3} under Regime 3—losses that are four to five times larger than the world-level costs and opposite in sign to the safe-bloc gains. The peripheral economy absorbs the concentrated architecture externality almost entirely. Under Regime 3, the pale red buffer bar partially offsets this

Figure 6: Architecture vs. Buffers: The Welfare Effects of the Balance Sheet Mechanism



Note: Three country perspectives. Solid bars: architecture cost (CE welfare of Regime 2 or 3 relative to Regime 4, buffers off). Pale bars: incremental buffer effect (activating $\phi_{\text{div}} = 0.5$, $\chi_E = 0.025$, $\zeta_E = 0.05$ within each regime). Left: world welfare ΔCE_W ($\times 10^{-3}$). Centre: safe-bloc average ΔCE_{HF} ($\times 10^{-4}$—note different scale). Right: R-country welfare ΔCE_R ($\times 10^{-3}$). The centre panel shows an unexpected positive solid bar for Regime 2: H and F households gain from concentrated H-denomination relative to the knife-edge, because the safe-haven inflow during a risk-off shock benefits H-currency stablecoin holders. The right panel concentrates the costs: R bears architecture losses of -4.6×10^{-3} under Regime 2 and -5.4×10^{-3} under Regime 3, recovered only partially by buffers under Regime 3. The architecture that benefits the safe bloc most is the same one that damages the periphery most. H and F country-specific decompositions available upon request.

with a recovery of approximately $+0.6 \times 10^{-3}$ through the val_{fx} channel, leaving a net R welfare loss of -4.8×10^{-3}—more than 80% of the architecture damage remains after buffers. Under Regime 2, the pale blue buffer bar is essentially zero: the val_{fx} channel does not exist under aligned backing, so buffers have no lever to pull for R.

Reading the centre and right panels side by side produces the distributional verdict: the architecture most beneficial to the safe bloc (Regime 2, positive centre bar) is the architecture most harmful to the periphery (Regime 2, deeply negative right bar). Capital regulation cannot resolve this conflict because it operates within-regime rather than across regimes. Only moving to the knife-edge (Regime 4) zeros out both bars simultaneously, eliminating both the safe-bloc gain and the R-country loss in one step.

4.6 Welfare Architecture and the New Triffin Dilemma

The analysis thus far characterizes the *channels* through which stablecoin design affects macroeconomic dynamics. To identify which channel is the dominant driver of *welfare*, we compute conditional consumption-equivalent (CE) welfare for each country across a fine grid of design parameter values. For each margin (D, A, U) and each bilateral edge of the currency simplex (HF, HR, FR), we report two statistics: the *range*—the difference between maximum and minimum CE welfare as the parameter traverses the edge—and the *argmax*—the edge value that maximizes welfare for each country and for the world. Table 2 summarizes the results.

The welfare hierarchy. The dominant result in Table 2 is the unit-of-account margin's supremacy. Along the HR edge—the edge that trades off safe-haven and risky-currency weight in the basket—the UoA welfare range is 6.04×10^{-3} in Regime 2, 6.05×10^{-3} in Regime 3, and 4.77×10^{-3} in Regime 4. The corresponding denomination ranges are 0.90, 0.20, and 0.03×10^{-3}, yielding U/D ratios of 6.7, 30.8, and 190.1 respectively. The backing margin lies between them at approximately 0.77×10^{-3} in Regimes 2–3 and 0.28×10^{-3} in Regime 4. The hierarchy is unambiguous: $\text{Range}^U \gg \text{Range}^A \gg \text{Range}^D$.

The mechanism is transparent. The denomination channel works through financial market frictions—portfolio-balance wedges—and reaches households indirectly, via exchange rate volatility and output gaps. The unit-of-account channel hits households directly in the utility function: when R depreciates and the stablecoin basket is H-tilted, stablecoin inflation in R numeraire falls below domestic CPI, so the stablecoin becomes cheap in R units exactly when domestic money is being debased. R households obtain a first-order welfare gain with no financial intermediation required. This is the formal account of real-world dollarization, now operating through the stablecoin's basket composition rather than currency substitution.

Table 2: Welfare Architecture: Importance and Distributional Conflict

Regime	Margin	Welfare Range ($\times 10^3$)				Argmax: HR edge		
		HF edge		HR edge		World	Safe	R
		ΔCE_W	ΔCE_R	ΔCE_W	ΔCE_R	x^*_W	x^*_{HF}	x^*_R
R2: Conc., aligned	D	0.52	0.88	0.90	2.26	1.00	0.88	1.00
	A	0.06	0.03	0.77	0.82	1.00	1.00	1.00
	U	**0.49**	**0.86**	**6.04**	**7.77**	**0.67**	**0.91**	**0.11**
R3: Conc., misaligned	D	0.20	0.37	0.20	0.74	0.47	0.12	0.71
	A	0.06	0.03	0.77	0.82	1.00	1.00	1.00
	U	**0.11**	**0.36**	**6.05**	**8.64**	**0.64**	**0.87**	**0.11**
R4: Knife-edge	*D*	*0.00*	*0.00*	*0.03*	*0.34*	—	—	—
	A	*0.00*	*0.00*	0.28	0.16	0.74	0.83	0.56
	U	***0.00***	***0.00***	**4.77**	**8.45**	**0.65**	**0.91**	**0.13**

Note: Welfare range (max minus min conditional CE welfare as the design parameter traverses the edge, $\times 10^3$) and welfare-maximising argmax for the HR edge, by regime and design margin. On the HR edge the parameter moves from 0 (full R weight) to 1 (full H weight). FR results are numerically symmetric to HR and omitted. Bold: unit-of-account margin. Italics: near-zero ($< 0.01 \times 10^3$).

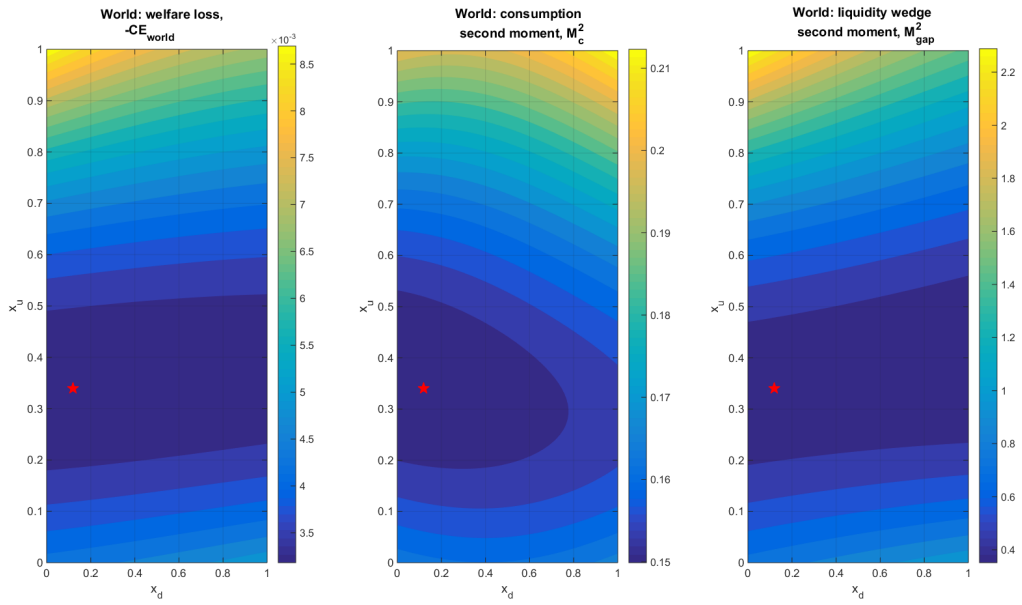
The knife-edge row in Table 2 provides a striking corollary. The denomination range on the HF edge collapses to 0.001×10^{-3}—numerically zero—while the UoA range on the HR edge remains 4.77×10^{-3}, fully 79 percent of its Regime 2 value. The knife-edge eliminates denomination-channel welfare effects completely while leaving the UoA welfare gradient entirely intact. *Denomination regulation and UoA governance are welfare-independent policy instruments.*

Figure 7 maps the joint variation of denomination and UoA across the full two-dimensional simplex edge, rather than slicing along one dimension at a time. The contour structure makes the hierarchy of Table 2 visually immediate.

In the left panel (welfare loss), the contours run nearly horizontal: they are lines of approximately constant x_u, so moving rightward along the x-axis—increasing the denomination R-weight x_d—crosses almost no contours, while moving upward—increasing the UoA R-weight x_u—crosses contours rapidly, tracking from the deep-blue welfare minimum at the bottom to the yellow maximum at the top. Denomination is the flat direction; unit of account is the steep direction.

The right panel (liquidity gap second moment M^2_{gap}) reproduces almost exactly the same contour geometry as the welfare loss panel. The near-perfect visual correlation identifies the mechanism: the UoA welfare gradient is inherited from the liquidity gap term in the utility

Figure 7: The Contour Structure and the Dominance of the Unit-of-Account Channel



Note: All three panels vary x_d (denomination R-weight, x-axis) and x_u (unit-of-account R-weight, y-axis) simultaneously along the safe-bloc to RoW ridge, with backing aligned to denomination throughout. *Left*: world welfare loss $-\text{CE}_{\text{world}}$—contours run nearly horizontally, confirming that welfare is steep in the UoA direction and flat in the denomination direction. *Centre*: world consumption second moment M_c^2—interior curvature reveals the denomination channel operating through consumption volatility; the oval around ($x_d \approx 0.6$, $x_u \approx 0.25$) is the denomination channel's footprint. *Right*: world liquidity gap second moment M^2_{gap}—contour structure mirrors the welfare loss panel almost exactly, establishing that the UoA welfare gradient is inherited from the liquidity gap channel in the utility function. The contrast between centre and right is itself a result: the denomination channel generates consumption volatility but is partially offset in general equilibrium welfare; the UoA channel transmits directly and without attenuation through the liquidity gap.

function. When the basket is R-heavy, the stablecoin provides purchasing-power insurance precisely when R households need it most—domestic money depreciates, but stablecoin real value does not—so the liquidity gap narrows and welfare rises.

The centre panel (consumption second moment M_c^2) tells a different story: it shows pronounced interior curvature, with an oval region around ($x_d \approx 0.6$, $x_u \approx 0.25$) that is absent from the other two panels. This is the denomination channel's footprint—it operates through consumption volatility but is partially offset in general equilibrium before reaching the utility function. The contrast between centre and right is itself a quantitative result: the denomination channel generates substantial consumption volatility but contributes only modestly to welfare loss, while the UoA channel passes through to welfare without attenuation.

The optimal UoA is interior. The world-welfare-maximising UoA composition is $x_W^* \approx 0.65$ across all active regimes—safe-tilted but not cornered. The optimal basket allocates roughly two-thirds of available simplex weight to the safe-currency bloc and one-third to R. This provides meaningful purchasing-power insurance to R households while containing the stablecoin demand volatility that destabilizes advanced-economy monetary transmission. Strikingly, x_W^* is stable at 0.64–0.67 regardless of denomination or backing: the optimal basket is a structural feature of the shock distribution and preference parameters, not of the architectural choice.

The New Triffin Dilemma. The argmax columns expose the distributional conflict. The advanced-economy safe bloc prefers $x_{HF}^* \approx 0.91$—near the corner that excludes R—because this minimizes stablecoin demand flows that destabilize their monetary transmission during R depreciation. R prefers $x_R^* \approx 0.12$—near the opposite corner, retaining substantial own-currency basket weight for maximum purchasing-power insurance. The world optimum $x_W^* \approx 0.65$ lies strictly between these positions and satisfies neither.

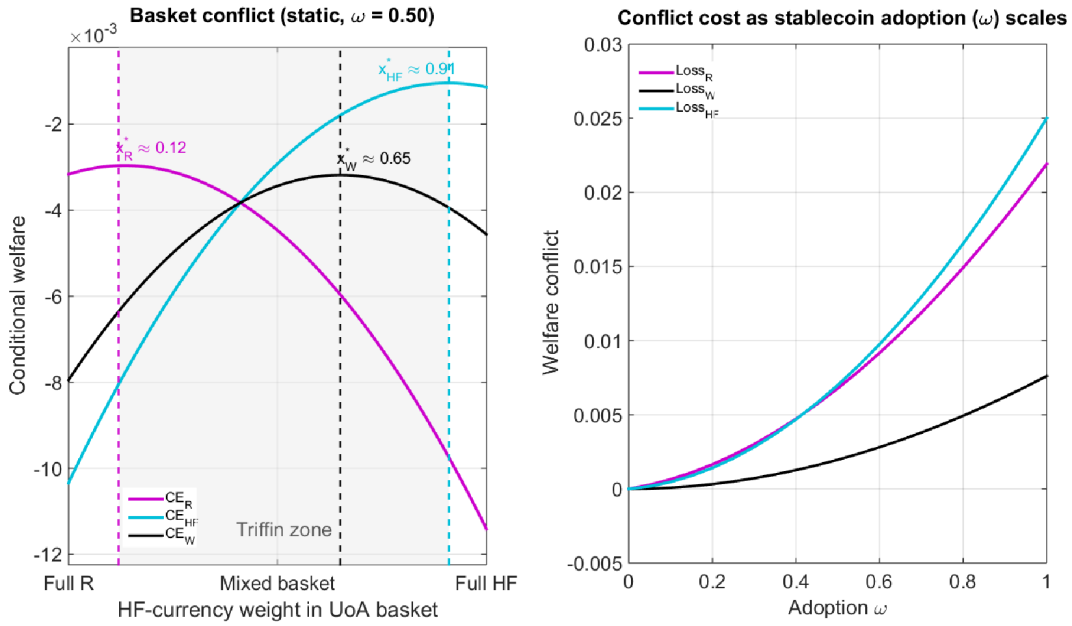
Proposition 5 (New Triffin Dilemma)**.** *Let x_j^* denote the welfare-maximising UoA edge parameter for country j on the HR simplex edge. Under the baseline calibration, $x_R^* < x_W^* < x_{HF}^*$ with $x_R^* \approx 0.12$, $x_W^* \approx 0.65$, and $x_{HF}^* \approx 0.91$. No single UoA vector simultaneously maximizes CE_R, CE_{HF}, and CE_W. The welfare gap between the R-optimal and the safe-bloc-optimal basket is increasing in adoption ω: the distributional conflict intensifies as stablecoin adoption grows.*

The conflict is not an artefact of the parameterization. It follows from the structure of Proposition 4: because the UoA channel operates through money demand without affecting UIP premia, the welfare cost to advanced economies of an R-inclusive basket comes entirely

through IS-curve stablecoin demand fluctuations. The larger R's weight in the basket, the greater the demand swings from R households during depreciation episodes, and the greater the IS disturbance transmitted to H and F through the stablecoin market. No denomination regulation can resolve this: the two instruments are orthogonal by construction.

Figure 8 renders Proposition 5 in its starkest form. The crossing of the CE_r and CE_{HF} curves is the dilemma in one image: moving the basket toward H (rightward) continuously improves advanced-economy welfare while reducing peripheral welfare from its interior optimum. There is no basket that avoids the trade-off; there is only a choice of where within the Triffin zone to locate. The world-welfare optimum sits near the midpoint of that zone.

Figure 8: The New Triffin Dilemma



Note: Welfare gradient along the HF-R unit-of-account edge. Left panel: static cross-section under Regime 4 (knife-edge denomination and backing) and stablecoin adoption fixed at its baseline ($\omega = 0.5$), with the UoA basket moving from full R-currency weight in the UoA basket (left x-axis) to full HF-safe-bloc weight (right x-axis). The crossing of CE_R (magenta) and CE_{HF} (cyan) illustrates the distributional conflict: no single basket simultaneously maximizes peripheral and advanced-economies core welfare. The world-welfare optimum $x^*_W \approx 0.65$ (vertical black line) lies strictly between the R optimum $x^*_R \approx 0.12$ (vertical magenta line) and the safe-bloc optimum $x^*_{HF} \approx 0.91$ (vertical cyan line). The shaded "Triffin zone" marks the set of baskets over which the distributional conflict is active. (Bilateral F-R and H-R edge results are numerically identical.) Right panel: Welfare conflict as a function of adoption (ω), using the fixed basket points $x^*_R = 0.12$, $x^*_W = 0.65$, $x^*_{HF} = 0.91$. As adoption rises, the welfare losses from moving toward the opposing bloc's preferred basket increase, making the distributional conflict quantitatively more important.

The comparison with the original Triffin dilemma is precise. Triffin observed that the Bretton Woods system required the US to run gold-eroding deficits in order to supply the

liquidity the world economy demanded. The two objectives—liquidity supply and gold credibility—shared the same instrument (the dollar-gold price) and therefore could not be simultaneously optimized. The new dilemma requires a global stablecoin issuer to provide purchasing-power insurance to the periphery (R-inclusive basket) in a way that destabilizes the core (demand volatility), and the two objectives again share the same instrument—the UoA basket. But crucially, the new dilemma has a dimension the original lacked: Proposition 4 guarantees that the financial stability instrument (denomination) and the distributional instrument (UoA) are orthogonal. Denomination regulation can be imposed without touching the distributional question. This separability is what makes the new dilemma tractable, even if the political economy of resolution is not.

Scale intensification and forward-looking urgency. The right panel of Figure 8 plots welfare conflict as a function of adoption $\omega \in [0, 1]$, evaluated at the fixed basket points $x_R^* = 0.12$, $x_W^* = 0.65$, and $x_{HF}^* = 0.91$. Three results are immediate. First, the welfare conflict between bloc preferences is monotonically increasing in ω: as adoption grows, the cost of satisfying one bloc's preferred basket at the expense of the other's rises proportionally. The dilemma is structural and it intensifies with scale, not a transitional feature of early-stage adoption. Second, the burden falls asymmetrically: the gap between the R-optimal and world-optimal baskets grows faster with ω than the gap between the HF-optimal and world-optimal baskets. The periphery has more at stake in basket governance than the advanced-economy core, yet it has less bargaining power in any multilateral negotiation—the same asymmetry that defined the original dilemma in the Bretton Woods context. Third, at the empirically current adoption share of approximately $\omega \approx 0.015$, all regime differences are quantitatively negligible: the welfare consequences of architecture are not yet material. They become significant around $\omega \approx 0.10$–0.20, a threshold reachable within a decade at plausible growth rates. Regulators who require observable welfare damage before acting will therefore be regulating after the conflict has already entrenched—precisely the forward-looking failure the original Triffin dilemma illustrated. Full results are reported in the Online Appendix (Figure OA.1).

5 Policy Implications

The five results above imply a regulatory framework built around two instruments, reflecting the orthogonality of the denomination and unit-of-account channels. Each instrument corresponds to a distinct institutional mandate; neither can substitute for the other.

Instrument 1: Denomination regulation for financial stability. The instrument requires symmetric diversified denomination ($\theta_j^d = 1/3$ for all j) as a financial-stability condition. This is the unique architecture that eliminates denomination-channel UIP distortions and cross-market contagion exactly (Propositions 1 and 2). The welfare cost of this requirement is negligible: the denomination range is at most 2.2×10^{-3} in CE welfare (Table 2), making it a near-Pareto improvement relative to concentrated alternatives. Denomination regulation belongs with financial stability authorities (FSB, BIS), who can set and enforce concentration limits on stablecoin liability composition in the same way they set large-exposure limits for banks.

Issuers who demonstrate symmetric denomination and denomination-aligned backing qualify for a regulatory safe harbour: simplified treatment, lower capital surcharges, access to central bank liquidity facilities. Issuers who deviate face macroprudential surcharges proportional to their concentration coefficients $|\theta_H^d - \theta_F^d|$, $|\theta_H^d - \theta_R^d|$, and misalignment norms $|\theta_j^a - \theta_j^d|$. Reserve quality requirements remain necessary but are not sufficient: Figure 1 shows that a stablecoin with 100 percent high-quality H-currency backing (Regime 2) generates identical H-F contagion to one holding risky R assets (Regime 3). Reserve quality and denomination structure are different regulatory objects addressing different risks.

Instrument 2: UoA governance for distributional balance. The instrument sets the UoA basket at the world-welfare optimum $x_W^* \approx 0.65$, which Table 2 shows is stable across all architectures. This requires multilateral negotiation: the advanced-economy bloc prefers $x_{HF}^* \approx 0.91$ and the periphery prefers $x_R^* \approx 0.12$. The world optimum lies between these positions and cannot be implemented by any single party acting unilaterally. UoA governance belongs with a multilateral institution representing both core and periphery interests—the IMF is the natural candidate, given its existing mandate over the international monetary system and its SDR basket methodology. The UoA governance problem is, in effect, the question of what the private-sector analogue of SDR basket weights should look like in a world where private global liquidity provision has become quantitatively significant.

Capital requirements complement but cannot replace denomination regulation. Figure 4 shows that the equity-feedback coefficient $\chi_E > 0$ is stabilising under the knife-edge but can amplify under concentrated denomination once the equity buffer is exhausted. The correct regulatory sequencing is: require knife-edge denomination first, then set capital requirements as a second line of defence for residual risks from backing misalignment and UoA volatility. Regulating capital without first requiring denomination diversification risks providing apparent stability in normal times followed by amplification during the stress

episodes when stability matters most.

Multipolar system perspective. The New Triffin Dilemma is not a temporary problem associated with early stablecoin adoption; it is structural, and it intensifies with scale. As the world transitions toward a multi-reserve configuration, the currency composition of the global stablecoin system is a first-order international monetary governance object—as consequential as reserve adequacy, and more consequential for welfare as demonstrated by Table 2. Dollar-denominated stablecoins reinforce the existing reserve hierarchy and concentrate denomination-channel risk in the dollar segment. The knife-edge design is neutral with respect to any reserve configuration, consistent with multipolarity in a way that concentrated-denomination designs are not. A well-governed stablecoin system—with denomination diversification enforced by financial stability authorities and basket composition negotiated multilaterally— could serve as a genuinely multipolar private liquidity mechanism. A poorly governed one will reproduce the Triffin dynamics at scale.

Beyond private stablecoin regulation, the orthogonality result raises a more ambitious institutional question. If denomination and unit-of-account governance can be separated into independent instruments, then a multilateral digital instrument designed from first principles—one that embeds knife-edge denomination at the protocol level and implements the world-welfare-optimal basket by multilateral mandate rather than private profit maximization—would resolve the New Triffin Dilemma by construction. The present paper establishes the theoretical conditions for such an instrument; its design lies outside our scope.

6 Conclusion

The immediate policy message of the paper is straightforward: regulation that focuses on reserve quality while leaving denomination architecture unregulated targets a second-order welfare lever. Stablecoin regulation that addresses denomination architecture without engaging with basket composition is solving the financial stability problem while leaving the dominant welfare problem—the New Triffin Dilemma—entirely unaddressed. A complete framework requires both instruments, delegated to separate institutions, operating on orthogonal channels.

The model underpinning this message has five structural results. Denomination diversification eliminates the contagion channel exactly: under the knife-edge, a peripheral risk-off shock that would normally spill into advanced-economy bilateral exchange rates through the stablecoin portfolio-balance mechanism leaves the H-F market completely undisturbed—at every shock realization, every period, every level of adoption (Propositions 1 and 2; Figure 1).

Backing alignment hedges issuer equity against exchange rate movements; misalignment creates a procyclical doom loop whose exhaustion dynamics reverse the sign of the equity-feedback stabilization (Proposition 3; Figure 4). The unit-of-account basket is structurally orthogonal to both, operating exclusively through domestic money demand (Proposition 4).

Welfare analysis reveals that orthogonality implies a hierarchy: the UoA channel generates welfare variation up to 190 times larger than the denomination channel (Table 2). The stablecoin functions primarily as insurance against peripheral currency risk. R households hold stablecoins not to exploit denomination arbitrage but to preserve purchasing power when their domestic currency depreciates. This is the formal account of dollarization—and it generates the New Triffin Dilemma (Proposition 5; Figure 8). The welfare-maximising UoA for the periphery is irreconcilably different from that of the advanced-economy core, and the conflict widens monotonically with adoption.

The separability result distinguishes the new dilemma from the original. Denomination regulation and basket governance address orthogonal channels and can be delegated to separate institutions. The original Triffin dilemma lacked this property—liquidity supply and dollar credibility used the same instrument—which is why it was insoluble within the Bretton Woods architecture. The new dilemma has an institutional solution, even if the political economy of achieving it across core and periphery interests is challenging.

Several extensions warrant future research. Endogenising the adoption decision would allow the model to characterize equilibria in which design choices feed back into adoption scale and hence into the magnitude of the Triffin conflict—important given that stablecoins are growing rapidly from their current 1.5 percent share of advanced- economy broad money. A political economy extension would formally characterize the private issuer's basket choice and the divergence between private and social optima that the New Triffin Proposition implies. A dynamic version—in which the distributional conflict evolves as the global monetary architecture shifts toward multipolarity—would connect the model to the longer-run question of whether private stablecoins can serve the role that SDR allocations were designed to serve, and at what governance cost.

The most consequential extension concerns not the private sector but the multilateral one. The separability result of Proposition 4 does more than identify two independent regulatory instruments: it identifies the precise conditions under which a *multilateral* digital liquidity instrument could resolve the New Triffin Dilemma from outside the private market. A multilateral issuer unconstrained by private profit maximization could implement the world-welfare-optimal basket x_W^* that no private issuer will voluntarily adopt. And a digitally native instrument—one whose denomination structure is enforced at the protocol level rather than through periodic political renegotiation—could implement knife-edge denomina-

tion continuously and automatically, eliminating the UIP distortions that arise whenever private issuers deviate from neutrality in response to competitive pressures. This combination—denomination neutrality by construction, welfare-optimal basket by multilateral mandate—is precisely what the current global financial safety net lacks. The existing toolkit (central bank swap lines, SDR allocations, IMF conditional programs) embodies structural trade-offs between speed, universality, conditionality, and political discretion that no architectural reform within the existing paradigm can fully resolve. Whether a digitally native multilateral liquidity instrument could serve as the missing governance layer—simultaneously reconciling speed with discipline, multilateral legitimacy with operational effectiveness, and crisis response with monetary sovereignty—is a question the present framework raises but does not answer. It is the natural next step in the research programme this paper initiates.

The broader implication is more pointed. A unipolar international monetary system operating within a multipolar world economy faces mounting strain. The rapid growth of private stablecoins intensifies rather than relieves it: without governance, private digital money reproduces the Triffin dynamics at scale and at speed. The orthogonality result in this paper suggests that digital infrastructure has a property the analogue monetary system does not: the ability to disentangle denomination, backing, and unit-of-account governance into three truly independent design dimensions, each addressable by the appropriate institution without contaminating the others. Whether that property can be harnessed to rewrite the institutional logic of the international monetary system—rather than merely to regulate a new class of private instrument—is the question that motivates the research agenda to which this paper contributes.

References

Tobias Adrian and Tommaso Mancini Griffoli. The rise of digital money. *IMF FinTech Notes*, 19(001), 2019. IMF FinTech Notes No. 19/001.

Raphael Auer, Giulio Cornelli, and Jon Frost. Rise of the central bank digital currencies: Drivers, approaches and technologies. BIS Working Paper 880, Bank for International Settlements, Basel, 2021.

Pierpaolo Benigno, Linda M. Schilling, and Harald Uhlig. Cryptocurrencies, currency competition, and the impossible trinity. *Journal of International Economics*, 141:103728, 2023.

Olivier Jean Blanchard and Charles M. Kahn. The solution of linear difference models under rational expectations. *Econometrica*, 48(5):1305–1311, 1980.

Markus K. Brunnermeier, Harold James, and Jean-Pierre Landau. The digitalization of money. NBER Working Paper 26300, National Bureau of Economic Research, Cambridge, MA, 2019.

Ricardo J. Caballero, Emmanuel Farhi, and Pierre-Olivier Gourinchas. An equilibrium model of "Global Imbalances" and low interest rates. *American Economic Review*, 98(1):358–393, 2008.

Ricardo J. Caballero, Emmanuel Farhi, and Pierre-Olivier Gourinchas. The safe assets shortage conundrum. *Journal of Economic Perspectives*, 31(3):29–46, 2017.

Guillermo A. Calvo. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383–398, 1983.

Christian Catalini, Alonso de Gortari, and Nihar Shah. Some simple economics of stablecoins. NBER Working Paper 30466, National Bureau of Economic Research, Cambridge, MA, 2022.

Richard Clarida, Jordi Galí, and Mark Gertler. A simple framework for international monetary policy analysis. *Journal of Monetary Economics*, 49(5):879–904, 2002.

Giancarlo Corsetti, Luca Dedola, and Sylvain Leduc. Optimal monetary policy in open economies. *Handbook of Monetary Economics*, 3:861–933, 2010.

Michael B. Devereux and Charles Engel. Monetary policy in the open economy revisited: Price setting and exchange-rate flexibility. *Review of Economic Studies*, 70(4):765–783, 2003.

Charles Engel. Currency misalignments and optimal monetary policy: A reexamination. *American Economic Review*, 101(6):2796–2822, 2011.

Emmanuel Farhi and Matteo Maggiori. A model of the international monetary system. *Quarterly Journal of Economics*, 133(1):295–355, 2018.

Xavier Gabaix and Matteo Maggiori. International liquidity and exchange rate dynamics. *Quarterly Journal of Economics*, 130(3):1369–1420, 2015.

Jordi Galí and Tommaso Monacelli. Monetary policy and exchange rate volatility in a small open economy. *Review of Economic Studies*, 72(3):707–734, 2005.

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

Gary Gorton and Jeffery Zhang. Taming wildcat stablecoins. Technical report, University of Chicago Law Review, 2023.

Pierre-Olivier Gourinchas and Hélène Rey. International financial adjustment. *Journal of Political Economy*, 115(4):665–703, 2007.

Pierre-Olivier Gourinchas, Hélène Rey, and Nicolas Govillot. Exorbitant privilege and exorbitant duty. Unpublished manuscript, University of California, Berkeley. Originally circulated as IMES Discussion Paper Series No. 2010-E-20, Bank of Japan, 2010, 2017.

Zhiguo He, Arvind Krishnamurthy, and Konstantin Milbradt. A model of the reserve asset. *American Economic Review*, 109(4):1230–1262, 2019.

Matteo Maggiori, Brent Neiman, and Jesse Schreger. International currencies and capital allocation. *Journal of Political Economy*, 128(6):2019–2066, 2020.

Tommaso Mancini Griffoli, Maria Soledad Martinez Peria, Itai Agur, Anil Ari, John Kiff, Adina Popescu, and Celine Rochon. Casting light on central bank digital currency. IMF Staff Discussion Note SDN/18/08, International Monetary Fund, Washington, DC, 2018.

Maurice Obstfeld and Kenneth Rogoff. Global implications of self-oriented national monetary rules. *Quarterly Journal of Economics*, 117(2):503–535, 2002.

Maurice Obstfeld, Jay C. Shambaugh, and Alan M. Taylor. Financial instability, reserves, and central bank swap lines in the panic of 2008. *American Economic Review: Papers and Proceedings*, 99(2):480–486, 2009.

Stephanie Schmitt-Grohé and Martín Uribe. Closing small open economy models. *Journal of International Economics*, 61(1):163–185, 2003.

Appendix

A The Log-Linearized Equilibrium System

This appendix presents the complete first-order log-linearized model, solved using Blanchard and Kahn (1980) conditions as implemented in Dynare 4.5 (Schmitt-Grohé and Uribe, 2003). All variables are log-deviations from the symmetric steady state. Exchange rate changes are denoted $\dot{e}_{jk,t} \equiv e_{jk,t} - e_{jk,t-1}$.

A.1 Supply. $\pi_{jd,t} = \beta \mathbb{E}_t \pi_{jd,t+1} + \kappa y_{j,t}$, $j \in \{H, F, R\}$.

A.2 Monetary policy. $i_{j,t} = \rho_i i_{j,t-1} + (1-\rho_i)(\phi_\pi \pi_{j,t} + \phi_y y_{j,t}) + \varepsilon^i_{j,t}$.

A.3 CPI inflation. Each country's CPI combines domestic and imported goods, with import share γ equally split across the two trading partners:

$$\begin{aligned}
\pi_{H,t} &= (1-\gamma)\pi_{Hd,t} + \tfrac{\gamma}{2}(\pi_{Fd,t} + \dot{e}_{hf,t}) + \tfrac{\gamma}{2}(\pi_{Rd,t} + \dot{e}_{hr,t});\\
\pi_{F,t} &= \tfrac{\gamma}{2}(\pi_{Hd,t} - \dot{e}_{hf,t}) + (1-\gamma)\pi_{Fd,t} + \tfrac{\gamma}{2}(\pi_{Rd,t} + \dot{e}_{fr,t});\\
\pi_{R,t} &= \tfrac{\gamma}{2}(\pi_{Hd,t} - \dot{e}_{hr,t}) + \tfrac{\gamma}{2}(\pi_{Fd,t} - \dot{e}_{fr,t}) + (1-\gamma)\pi_{Rd,t}.
\end{aligned}$$

A.4 Terms of trade. Bilateral terms of trade evolve as:

$$\begin{aligned}
\dot{\text{tot}}_{hf,t} &= \dot{e}_{hf,t} + \pi_{Fd,t} - \pi_{Hd,t};\\
\dot{\text{tot}}_{hr,t} &= \dot{e}_{hr,t} + \pi_{Rd,t} - \pi_{Hd,t};\\
\dot{\text{tot}}_{fr,t} &= \dot{e}_{fr,t} + \pi_{Rd,t} - \pi_{Fd,t}.
\end{aligned}$$

Only two are independent: $\text{tot}_{fr,t} = \text{tot}_{hr,t} - \text{tot}_{hf,t}$ by construction.

A.5 Stablecoin basket inflation. $\pi^{\text{sc}}_t = \theta^u_H \pi_{Hd,t} + \theta^u_F(\pi_{Fd,t} + \dot{e}_{hf,t}) + \theta^u_R(\pi_{Rd,t} + \dot{e}_{hr,t})$.

A.6–A.9 MB-SC demand wedges. $mb_{j,t} = -\eta i_{j,t}$; $mb_{H,t} - s_{H,t} = \eta(\pi^{\text{sc}}_t - \pi_{H,t})$; $mb_{F,t} - s_{F,t} = \eta(\pi^{\text{sc}}_t - \dot{e}_{hf,t} - \pi_{F,t})$; $mb_{R,t} - s_{R,t} = \eta(\pi^{\text{sc}}_t - \dot{e}_{hr,t} - \pi_{R,t})$.

A.10–A.11 Liquidity services. $l_{j,t} = (1-\omega)mb_{j,t} + \omega s_{j,t}$; $\text{gap}_{j,t} = c_{j,t} - l_{j,t}$.

A.12 Euler equations. $(\sigma + \varphi_l)c_{j,t} - \varphi_l l_{j,t} = \mathbb{E}_t[(\sigma + \varphi_l)c_{j,t+1} - \varphi_l l_{j,t+1}] - (i_{j,t} - \mathbb{E}_t \pi_{j,t+1})$.

A.13–A.15 UIP and exchange rates.

$$
\begin{aligned}
e_{hf,t} &= \mathbb{E}_t e_{hf,t+1} + i_{H,t} - i_{F,t} - \psi_{hf,t};\\
e_{hr,t} &= \mathbb{E}_t e_{hr,t+1} + i_{H,t} - i_{R,t} - \psi_{hr,t};\\
e_{fr,t} &= \mathbb{E}_t e_{fr,t+1} + i_{F,t} - i_{R,t} - \psi_{fr,t}.
\end{aligned}
$$

A.16–A.21 Risk premia.

$$
\begin{aligned}
\psi^{\text{sc}}_{hf,t} &= \psi_b(\theta^d_H - \theta^d_F)\bar{s}(s_t - \zeta_E E^{\text{sc}}_t);\\
\psi^{\text{sc}}_{hr,t} &= \psi_b(\theta^d_H - \theta^d_R)\bar{s}(s_t - \zeta_E E^{\text{sc}}_t);\\
\psi^{\text{sc}}_{fr,t} &= \psi_b(\theta^d_F - \theta^d_R)\bar{s}(s_t - \zeta_E E^{\text{sc}}_t);\\
\psi_{hf,t} &= \psi_b b_{hf,t} + \psi^{\text{sc}}_{hf,t} - \chi_E E^{\text{sc}}_t;\\
\psi_{hr,t} &= \psi_b b_{hr,t} + \psi^{\text{sc}}_{hr,t} - \chi_E E^{\text{sc}}_t + \nu_t;\\
\psi_{fr,t} &= \psi_b b_{fr,t} + \psi^{\text{sc}}_{fr,t} - \chi_E E^{\text{sc}}_t + \nu_t.
\end{aligned}
$$

A.22–A.26 Issuer block.

$$
\begin{aligned}
r^A_t &= \theta^a_H i_{H,t} + \theta^a_F(i_{F,t} + \mathbb{E}_t \dot{e}_{hf,t+1}) + \theta^a_R(i_{R,t} + \mathbb{E}_t \dot{e}_{hr,t+1});\\
\pi_{\text{prof},t} &= \bar{s}(r^A_t - \pi^{\text{sc}}_t);\\
\text{val}_{q,t} &= \bar{s}\sum_j \theta^a_j \Delta q_{j,t};\\
\text{val}_{fx,t} &= \bar{s}[(\theta^a_F - \theta^d_F)\dot{e}_{hf,t} + (\theta^a_R - \theta^d_R)\dot{e}_{hr,t}];\\
E^{\text{sc}}_t &= (\beta^{-1} - \phi_{\text{div}})E^{\text{sc}}_{t-1} + \pi_{\text{prof},t} + \text{val}_{q,t} + \text{val}_{fx,t}.
\end{aligned}
$$

A.27–A.32 External block and goods markets. $b_{jk,t} = \beta^{-1} b_{jk,t-1} + nx_{jk,t}$; $nx_{jk,t} = \frac{\gamma}{2}(c_{k,t} - c_{j,t}) + \frac{\gamma}{2}\,\text{tot}_{jk,t}$; $y_{j,t} = c_{j,t} + nx_{j,t} - T_{j,t}$.

B Proofs

Proof of Proposition 1. Substituting $\theta^d_H = \theta^d_F = \theta^d_R = 1/3$ into (A.16)–(A.18): each concentration coefficient $(1/3 - 1/3) = 0$, so $\psi^{\text{sc}}_{hf,t} = \psi^{\text{sc}}_{hr,t} = \psi^{\text{sc}}_{fr,t} = 0$ for all t, all s_t, E^{sc}_t, and all shock realizations. The result holds for any ψ_b, $\bar{s}$, ζ_E, ω. □ □

Proof of Proposition 2. Under $\theta^d_H = 1$, $\theta^d_F = \theta^d_R = 0$: $\psi^{\text{sc}}_{hf,t} = \psi_b \cdot \bar{s} \cdot s_t$. Any $\Delta s_t > 0$ induced by H-R/F-R shocks enters $\psi^{\text{sc}}_{hf,t}$ with positive weight, shifting $e_{hf,t}$ through (A.13). Under $\theta^d_H = 1/3$, $\psi^{\text{sc}}_{hf,t} = 0$ identically by Proposition 1. □ □

Proof of Proposition 3. From (A.25): $\text{val}_{fx,t} = \bar{s}[(\theta_F^a - \theta_F^d)\dot{e}_{hf,t} + (\theta_R^a - \theta_R^d)\dot{e}_{hr,t}]$. Under alignment $\theta_j^a = \theta_j^d$ for all j, so all coefficients vanish and $\text{val}_{fx,t} = 0$ for every exchange rate path. Under misalignment with $\theta_R^a > \theta_R^d$ and $\dot{e}_{hr,t} < 0$ (R depreciates): $(\theta_R^a - \theta_R^d)\dot{e}_{hr,t} < 0$, giving $\text{val}_{fx,t} < 0$: equity is destroyed procyclically. □ □

Proof of Proposition 4. θ^u enters exclusively through π_t^{sc} in (A.5). Tracing downstream: π_t^{sc} appears in (A.7)–(A.9) and (A.23) only. It does not appear in (A.16)–(A.21) or (A.25). The indirect path $\theta^u \to \pi_t^{\text{sc}} \to (mb_j - s_j) \to s_t \to \psi_{jk,t}^{\text{sc}}$ has slope $\psi_b(\theta_j^d - \theta_k^d)\bar{s}$, determined by θ^d not θ^u. At the knife-edge this is zero by Proposition 1; in general, UoA and denomination channels are structurally orthogonal. □ □

C Blanchard–Kahn Conditions

The system is written as $\mathbf{A}\mathbb{E}_t\mathbf{x}_{t+1} = \mathbf{B}\mathbf{x}_t + \mathbf{C}\boldsymbol{\varepsilon}_t$. The model has **fifteen** forward-looking (jump) variables: three consumption variables $\{c_H, c_F, c_R\}$ and three liquidity-service variables $\{l_H, l_F, l_R\}$ from the Euler equations (A.12); three bilateral exchange rates $\{e_{hf}, e_{hr}, e_{fr}\}$ from the UIP conditions (A.13)–(A.15); three domestic inflation rates $\{\pi_{Hd}, \pi_{Fd}, \pi_{Rd}\}$ from the Phillips curves (A.1); and three CPI inflation rates $\{\pi_H, \pi_F, \pi_R\}$ appearing in the Euler equations through $\mathbb{E}_t\pi_{j,t+1}$ and in monetary policy (A.2). Issuer equity E^{sc} and bilateral net foreign asset positions b_{jk} are predetermined state variables.

Dynare's `check` command confirms exactly **fifteen** unstable eigenvalues of $\mathbf{A}^{-1}\mathbf{B}$ across all four design regimes at the baseline calibration, satisfying the Blanchard–Kahn order condition for determinacy.

NFA stationarity follows from $\psi_b > 0$: the portfolio-balance premium responds to bilateral NFA deviations, providing the standard stationarity device of Schmitt-Grohé and Uribe (2003). Under the knife-edge, the stablecoin block decouples from the UIP block at first order by Proposition 1. The eigenvalue count is unchanged; the impulse responses of e_{hf} and ψ_{jk}^{sc} to the ν shock are exactly zero by construction.